arXiv:1404.4812v2 [quant-ph] 10 Aug 2015

Beyond Bell's Theorem II: Scenarios with arbitrary causal structure

Abstract. It has recently been found that Bell scenarios are only a small subclass of interesting setups for studying the non-classical features of quantum theory within spacetime. We find that it is possible to talk about classical correlations, quantum correlations and other kinds of correlations on any directed acyclic graph, and this captures various extensions of Bell scenarios which have been considered in the literature. From a conceptual point of view, the main feature of our approach is its high level of unification: while the notions of source, choice of setting and measurement play all seemingly different roles in a Bell scenario, our formalism shows that they are all instances of the same concept of "event".

Our work can also be understood as a contribution to the subject of causal inference with latent variables. Among other things, we introduce hidden Bayesian networks as a generalization of hidden Markov models.

Contents

1. Introduction 2
2. What is causal structure? 8
3. Correlations 10
4. Classical correlations 13
5. Correlations within any theoretical framework 22
6. Quantum correlations 31
7. Classical correlations as hidden Bayesian networks 36
Appendix A. Measure-theoretic technicalities 40
Appendix B. Finite approximations of probability spaces 43
References 45

2010 *Mathematics Subject Classification.* Primary: 81P05, 62A01; Secondary: 28A35.

Key words and phrases. Bell's theorem, quantum nonlocality, causal inference, hidden Bayesian networks, measure-theoretic probability theory.

Acknowledgements. Thanks to all the participants of the Benasque meeting on Causal Structure in Quantum Theory, and in particular to the organizers Lluís Masanes and Rob Spekkens. Matthew Pusey has kindly made many helpful suggestions and contributed to some of our proofs. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Economic Development and Innovation. The author has been supported by the John Templeton Foundation.

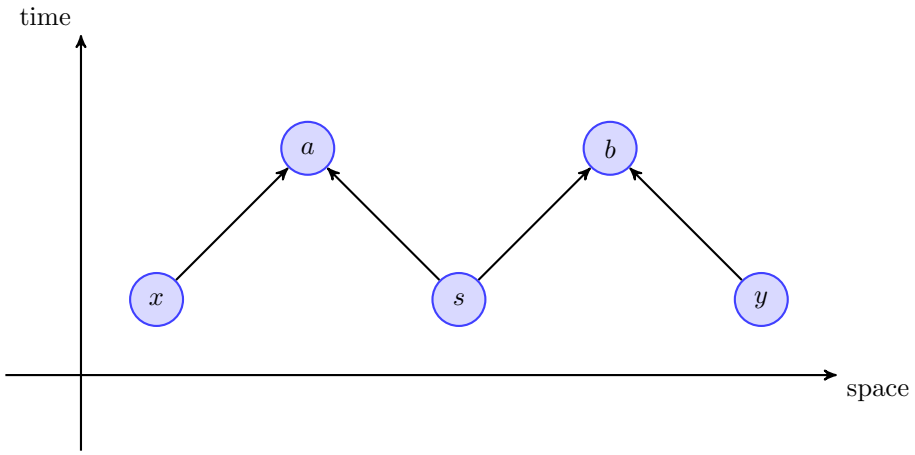

FIGURE 1. Causal structure of the two-party Bell scenario [1].

1. Introduction

Bell's discovery of "quantum nonlocality" [1, 2] has maybe been one of the most influential discoveries in theoretical physics over the past 50 years. Not only has it been a vexing challenge and a source of philosophical discussion for many physicists interested in the foundations of quantum theory, it has also led more recently to the development of new information processing protocols whose performance and security rely crucially on violations of Bell inequalities [3, 4, 5, 6].

Over the course of time, it has been realized that Bell's *gedankenexperiment*, whose causal structure is illustrated in Figure 1, is only one of many possible setups in which the non-classicality of quantum theory manifests itself in correlations between between different observations distributed in spacetime. Generalizations of Bell's original idea include Svetlichny's extension from a two-party scenario to a many-party scenario [7], Popescu's "hidden nonlocality" [8] which modifies the Bell scenario by including additional preselection measurements (Figure 2), and the "bilocality" scenario of Branciard, Gisin and Pironio [9] (Figure 3). Subsequently, we proposed a general definition of what constitutes a *correlation scenario*, showed how this definition comprises two-party and many-party Bell scenarios, the bilocality scenario, and many new ones (like Figure 4), about which we also derived some first results [10]. Most recently, Gallego, Würflinger, Chaves, Acín and Navascués have considered "sequential correlation scenarios" [11] which extend Popescu's idea by adding an additional choice of measurement setting (Figure 5).

The purpose of this present paper is threefold:

(a) We propose a general definition of a class of *gedankenexperiments*, in which one can talk about classical correlations, quantum correlations and the like, and which comprises all the above scenarios as particular instances. We provide definitions of classical and quantum correlations and show that they correspond to the usual ones in the case of a Bell scenario. Moreover, we prove some very general technical results about classical correlations.

(b) We point out that our definition achieves a new level of unification: there is no distinction whatsoever between a "source", a "choice of measurement setting" and a "measurement". All three concepts are subsumed by our notion of **event**, as illustrated in the Bell scenario of Figure 1. A pair of events may or may not be connected by a causal link drawn as

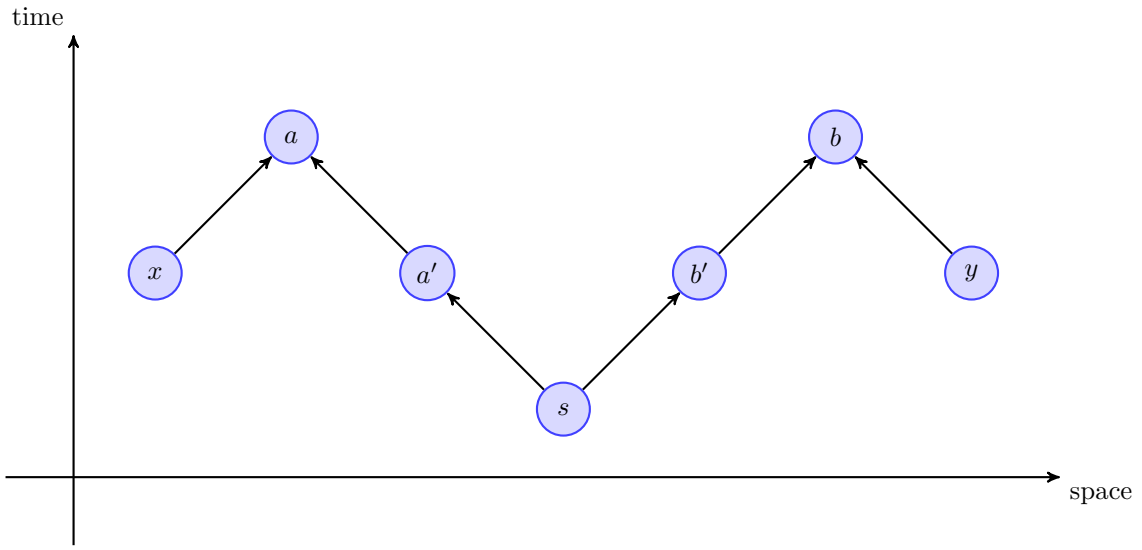

FIGURE 2. Causal structure of Popescu's scenario [8].

an arrow pointing from one event to the other. We think of an event as a measurement apparatus located at a specific point in spacetime which probes and potentially modifies the outside world and outputs a classical outcome. There are two ways to interpret the causal links: one can either think of them as representing a theoretical hypothesis about which events influence which other ones, or alternatively as physical systems propagating between the events—exactly as in Hardy's **operator-tensor formulation** of quantum theory [12]. Both of these interpretations are equally viable, since the underlying mathematics is the same.

(c) The first interpretation connects nicely with the third purpose of this paper: our definitions and results can also be understood as a contribution to the subject of **causal inference** [13]. Typically, in causal inference one starts with a joint probability distribution of a finite number of random variables and then asks which underlying causal structures, describing which variables influence which other ones, may have generated the given joint distribution. With our definitions, the notion of **hidden variable** or **latent variable** [13] becomes central to causal inference. The basic idea is as follows: let us assume that the joint probability distribution that we start with represents the distributions of outcomes of probings of a physical system at various points in spacetime, as outlined above. At each of these points, we sample the system, but we do not know whether the information that we obtain by that sampling is a complete representation of the state of the system at that particular point in spacetime. In other words, it may be that there is some "hidden information" propagating through that point in spacetime, while not being revealed by the probing. This seems like a reasonable assumption to make when dealing with systems for which we are uncertain whether our theoretical description is complete. We imagine that this applies in particular to biological systems, whose functioning can rarely be comprehensively understood; but it also applies to fundamental physics, as exemplified

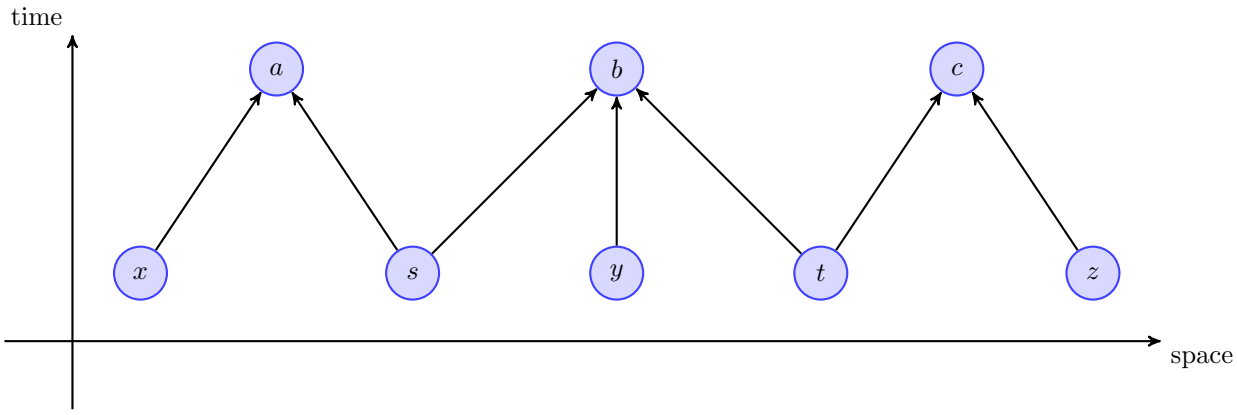

FIGURE 3. Causal structure of Branciard–Gisin–Pironio's "bilocality" scenario [9].

by Einstein's conviction that quantum mechanics was "incomplete" and hidden variables had to be introduced which would describe additional degrees of freedom that were hitherto unknown [14]. To summarize, we believe that any reasonable generalization of Bell scenarios needs to consider hidden variables not only behind events representing sources of (entangled) particles, but behind *every* spacetime event relevant to the gedankenexperiment. For the standard Bell scenario of Figure 1 and its many-party extensions, we will see in Propositions 3.3, 4.3 and 6.5 that the resulting formalism is equivalent to the conventional one.

Connections between causal inference and Bell's theorem have previously been pointed out by Wood and Spekkens [15]. Similar connections have also been explored by Henson, Lal and Pusey [16] and were also hinted at in our previous paper [10]. Most recently, they are being actively investigated by Chaves, Gross and collaborators [17, 18, 19, 20].

Summary and statement of results. Now that the main ideas have been explained, we outline in more detail our definitions and results, the proofs of which will sometimes constitute a technical tour de force. We suspect that some of our results may seem quite intuitive and unsurprising, but, as the technical difficulty of the proofs will demonstrate, they are non-trivial. In fact, despite sustained efforts, we have not achieved perfect rigor in all our arguments; the corresponding arguments are marked as "SKETCH OF PROOF" rather than "PROOF". We challenge any skeptical reader to find simpler and completely rigorous proofs.

We start out in Section 2 by discussing the relevant notion of causal structure, which is given by the concept of **directed acyclic graph**, just as in the subject of Bayesian networks and causal inference [13]. This graph describes which events have a *direct* influence on which other ones through causal links. During the course of our investigations, it will turn out that it is sufficient to know which events have a direct *or indirect* influence on which other ones, and it is not necessary to distinguish between a direct and an indirect causal link. In fact, Theorem 5.3 shows that only the **partial order** (or equivalently, **causal set**) induced by the graph needs to be known in order to discuss classical, quantum and other correlations on the given causal structure.

Section 3 continues by introducing the notion of **correlation** on a given causal structure. Roughly, a correlation is a joint distribution over the outcome variables situated at the events

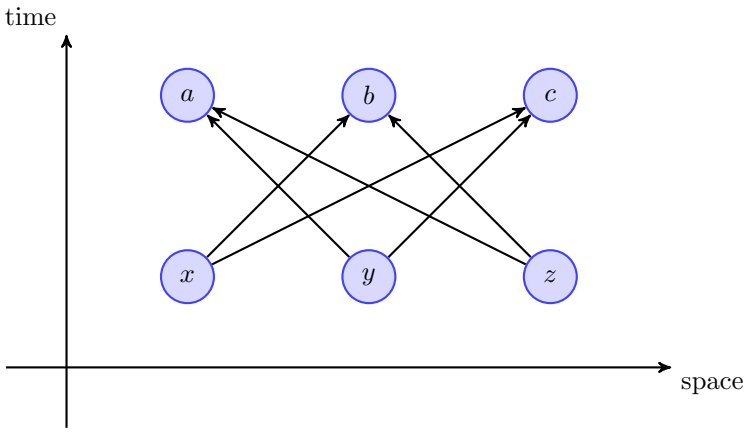

FIGURE 4. The triangle scenario from [10] in our present definition. In this two-dimensional drawing, the causal structure may look contrived; but in a three- or four-dimensional spacetime, the events can be arranged in such a way that the causal structure is indeed the one induced from the lightcones, for example by spatially arranging all six events as the vertices of a regular hexagon, while x, y and z happen at a suitably earlier time than a, b and c.

which makes two variables independent as soon as they have disjoint causal past. This formalizes Reichenbach's **common cause principle** "no correlation without causation" [21] in our framework. On a given causal structure, these independence conditions constitute a minimal set of requirements that one would expect an observed joint probability distribution to satisfy, assuming that it arises from that causal structure. This is analogous to the no-signalling condition in a Bell scenario, which is likewise a minimal requirement for a conditional distribution $P(ab|xy)$ to be regarded as potentially physical. As shown in Proposition 3.5, our definition essentially reduces to the no-signalling assumption in the case of a Bell scenario.

Then, in Section 4, we introduce hidden variables. A given correlation, as in the previous paragraph, is a **classical correlation** if it can be explained in terms of hidden variables being passed along the causal links and being processed at the nodes. This formalizes the above idea of causal inference with hidden variables: one asks which observational results are compatible with the given causal structure, while permitting unobserved information to propagate along the causal links. A given joint distribution is compatible with the causal structure in this sense if and only if it is a classical correlation. It is an essential feature of our definitions and results that we do not restrict a hidden variable to have a finite domain, but allow for it to live on an arbitrary measurable space. The measure-theoretical background necessary for dealing with hidden variables with arbitrary domain is outlined in Appendix A. We show in Proposition 4.2 that every classical correlation is in fact a correlation. Our further results, in parts quite difficult to prove, are the following: Proposition 4.3 characterizes classical correlations in a Bell scenario as those whose associated conditional distribution is "local" in the conventional sense [2]. Theorem 4.4 then shows that any classical correlation on any causal structure can be approximated arbitrarily well by other classical correlations for which the hidden variables take on only finitely many values; the proof relies on some auxiliary measure-theoretical results derived in Appendix B. Towards the end of

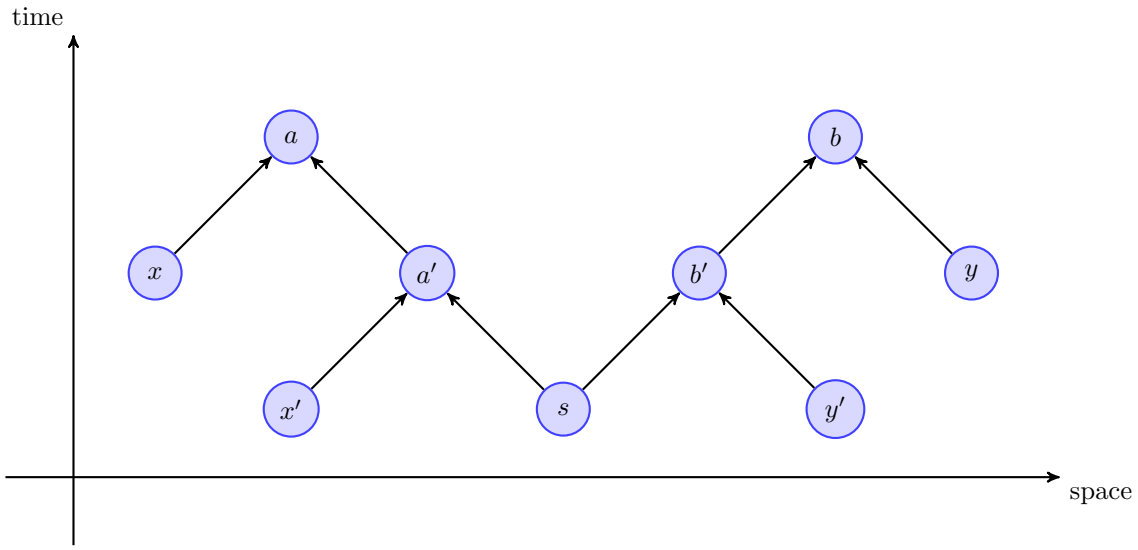

FIGURE 5. Causal structure of the two-party sequential correlation scenario with sequences of two measurements [11].

the section, Conjecture 4.5, together with its partial proof, concerns the question whether one can assume without loss of generality that the information processing from incoming to outgoing hidden variables at each event with incoming hidden variables is deterministic.

Motivated by the question of how to define quantum correlations within our formalism, Section 5 discusses the general properties that a theoretical framework must satisfy in order for a corresponding notion of correlation within that framework to exist. We find that, if a framework for physical theories is given by a symmetric monoidal category $\mathtt{C}$ satisfying suitable additional conditions, then there exists a notion of "$\mathtt{C}$-correlation". We explain how classical correlation are an instance of this general construction. In complete generality, Theorem 5.3 shows that the set of $\mathtt{C}$-correlations only depends on the causal structure as a partially ordered set instead of all the details of a directed graph.

Building on these considerations, we define quantum correlations in Section 6. We hypothesize in Conjecture 6.4 that every classical correlation is a quantum correlation, but due to the technical challenges involved, we can only offer a partial proof. Then, Proposition 6.5 shows how our definition of quantum correlation corresponds to the conventional one in the case of Bell scenarios.

Finally, Section 7 discusses another point of view on classical correlations. By analogy with hidden Markov models, we introduce hidden Bayesian networks. The main result here is Theorem 7.3, which shows that a given joint probability distribution over the outcome variables is a classical correlation if and only if it can be modelled by a hidden Bayesian network.

Terminology and Notation. We collect here the relevant notation that we use for graphs $G = (V, E)$. All the graphs G that we consider are *directed acyclic graphs* as in Definition 2.2, and hence we usually omit the qualifiers "directed acyclic" and simply refer to G as a *graph*.

While we always think of the nodes $v \in V$ as events in spacetime, we mostly call them "nodes" rather than "events", and regard these two terms as essentially synonymous. Similarly, we usually

speak of "edges" $e \in E$ rather than "causal links", although we think of the edges as causal links along which physical systems propagate. (With the exception of the edges of the hidden Bayesian networks of Section 7, which we do not interpret like this.)

For nodes $v \in V$ and edges $e \in E$, we put

$$\begin{aligned}
\mathrm{src}(e) &:= \text{source node of } e \\
\mathrm{tar}(e) &:= \text{target node of } e \\
\mathrm{pa}(v) &:= \text{parent nodes of } v &&= \{u \in V \mid u \to v\} \\
\mathrm{ch}(v) &:= \text{children nodes of } v &&= \{w \in V \mid v \to w\} \\
\mathrm{pst}(v) &:= \text{causal past of } v &&= \{u \in V \mid u \to \ldots \to v\}
\end{aligned}$$

We can also apply each of these functions to sets of edges (respectively nodes), in which case we put $\mathrm{pa}(U) := \bigcup_{v \in U} \mathrm{pa}(v)$, and similarly for the other functions. See Section 2 for more detail on all of this.

In order to keep the notation from being too confusing, we distinguish between a node and the corresponding outcome variable it carries. If $v \in V$ is a node, then the outcome variable at v is denoted $o[v]$, and $P(o[v])$ is its distribution. This notation is unfortunately somewhat unusual, since in the Bell scenario of Figure 1 we need to denote e.g. the outcome of Alice by "$o[a]$" instead of the conventional "a", but it will help to keep things clear.

Disclaimer. The definitions proposed in the present paper are not intended to be final or conclusive in any sense. Rather, they constitute only one possible approach for generalizing Bell scenarios to arbitrary causal structures, and we expect that other definitions are possible and potentially useful [16]. It may also be desirable to have definitions which even allow for the consideration of indefinite causal structure [22], although we do not know how to achieve this. Further investigations will have to tell which definitions are the most adequate and useful; the only thing that seems certain at the present time is that Bell scenarios constitute only a small subclass of all conceivable scenarios, and that there are many other interesting ones that one can consider [10, 16].

A potentially problematic aspect of our approach is this: in order to observe the joint probability distribution of outcomes of probings in an experimental situation, one needs to gather statistics, i.e. repeat the experiment many times *with independent trials*. This independence assumption of trials goes somewhat against the spirit of allowing hidden variables wherever conceivable, and might be hard to justify in certain applications, e.g. in neurophysiology, where the neurons may also store hidden information between the trials. At present, we do not know under which conditions this independence assumption can be relaxed, but the results of Bierhorst [23] for the case of the Bell scenario are encouraging.

2. What is causal structure?

Our goal is to study correlations between different observations conducted on a collection of systems and to determine which causal structures between these observations are compatible with these correlations. So we first need to discuss what the term "causal structure" actually means. Generally, a causal structure is defined on a collection of events V, and the function of a causal structure is to clarify which events may **influence** which other ones. Whenever event v may influence event w, we write $v \to w$. Upon taking the events to be nodes and regarding the influence relations $v \to w$ as edges, we therefore obtain a **directed graph** $G = (V, E)$, i.e. a set of nodes V together with a set of directed edges $E \subseteq V \times V$, which becomes the mathematical model of a causal structure. If there is no causal link from v to w, then we write $v \not\to w$. We will always assume that V is finite, and then so is E; it is probably possible to relax this assumption, albeit at the cost of additional technical challenges. Any edge $e \in E$ is necessarily of the form $v \to w$ for a unique pair of nodes $v, w \in V$, which we call the **source** and **target** of e and also denote by $\mathrm{src}(e)$ and $\mathrm{tar}(e)$. For $v \in V$, we write

$$\mathrm{in}(v) := \{\, e \in E \mid \mathrm{tar}(e) = v \,\}$$

for the collection of **incoming edges** at v, and similarly

$$\mathrm{out}(v) := \{\, e \in E \mid \mathrm{src}(e) = v \,\}$$

for the **outgoing edges** at v. The upcoming acyclicity condition of Definition 2.2 will guarantee that these two sets are disjoint. In Section 7, we will use similar concepts for nodes: the set

$$\mathrm{pa}(v) := \{\, u \in V \mid u \to v \,\}$$

is the collection of **parents** of v, while

$$\mathrm{ch}(v) := \{\, w \in V \mid v \to w \,\}$$

are its **children**.

We typically think of the events or nodes of G as points in spacetime, and there is an edge or causal link $v \to w$ if w lies in the future lightcone of v. However, our mathematical results are completely general and do not depend on this interpretation. In particular, we are confident that other ways for obtaining candidate causal structures are feasible and interesting, in particular in applications to mathematical modelling in other fields like biology or climate science. In these situations, it will often be the case that it is not known which events can influence which other ones, and therefore the causal structure is not known in advance. Applying our formalism to these cases is nevertheless possible by first choosing some "trial" causal structures and then asking, for each of these, whether it is compatible with the given observations or not. We do not subscribe to any particular method for finding such trial causal structures; doing so is up to the user of our formalism. It should be an interesting challenge to come up with algorithms for doing so, possibly in analogy with the standard causal inference algorithms commonly used [13, p. 50ff].

In general, even if $u \not\to w$, it may be possible that there is an **indirect** causal link from u to w proceeding through other nodes:

2.1. Definition. *A **directed path** from $u \in V$ to $w \in V$ is a finite sequence of nodes $v_1, \ldots, v_n \in V$ with $u \to v_1 \to \cdots \to v_n \to w$.*

We use the shorthand $u \rightsquigarrow w$ for the existence of a directed path from u to w, and $u \not\rightsquigarrow w$ for the absence of any such directed path. In the conventional formalism of Bayesian networks [13], no information can propagate along indirect causal links since the intermediate nodes "screen off" any

potential indirect influence of u on w, and hence the outcome variable $o[u]$ may become conditionally independent of $o[w]$ given $o[v]$ (whether this happens depends on the overall structure of the graph). For us, this is no longer the case: the outcome variable $o[v]$ living on an intermediate node v constitutes merely a probing of the environment without the guarantee that it faithfully represents all information propagating through v, and hence we cannot expect $o[v]$ to screen off any correlations between $o[u]$ and $o[w]$.

The additional property enjoyed by "$\rightsquigarrow$" is **transitivity**:

$$u \rightsquigarrow v, \quad v \rightsquigarrow w \quad \Longrightarrow \quad u \rightsquigarrow w.$$

This makes sense if one interprets "$u \rightsquigarrow v$" as "u is in the causal past of v". So the information carried by $\rightsquigarrow$ is exactly that of a **partially ordered set**, which for us plays a rôle similar to that of a **causal set** [24]. Due to the frequent occurrence of the concept, we introduce a special notation for the causal past,

$$\mathrm{pst}(v) := \{\, v \,\} \cup \{\; u \in V \;\mid\; u \rightsquigarrow v \;\}\,.$$

It is important to note that we consider every node itself to be a member of its causal past.

A condition that any reasonable causal structure should satisfy is that there are no causal loops, i.e. that no event can influence itself in any non-trivial way:

2.2. Definition (e.g. [13]). *A **directed acyclic graph** is a directed graph with $v \not\rightsquigarrow v$ for all $v \in V$.*

Note that this definition is sensible only if we do not allow paths of length zero in Definition 2.1.

For example, Figures 1–5 depict directed acyclic graphs. Acyclicity means that there is no directed sequence of edges which returns to its starting point. In particular, no node has a self-loop.

The particular causal structure considered by Bell [1] is the one of Figure 1. It is worth noting that the five nodes of that figure are normally considered to play very different roles: s represents a **source** emitting particle pairs or other physical systems, one in each direction. It is usually not considered to yield an outcome $o[s]$, although it almost certainly does in any actual experiment, where the outcome $o[s]$ in any given run will be one of "successfully produced an entangled particle pair" or "failed to produce an entangled particle pair". The nodes x and y stand for a **choice of measurement** randomly selected by an agent (or by an apparatus, in most actual experiments) with values $o[a]$ and $o[b]$. Finally, the events a and b represent the execution of these chosen measurements on the systems emitted by the source which produce outcomes $o[a]$ and $o[b]$. However, this categorization of the events into three different kinds is completely superfluous: one can consider all five variables $o[s]$, $o[x]$, $o[y]$, $o[a]$, $o[b]$ to be outcomes of observations, which are (probabilistic) functions of the values of the incoming hidden variables propagating along the incoming causal links (if any), and emitting hidden variables along the outgoing causal links (if any). Here, the "hidden variables" may be classical, quantum or other kinds of systems; we will show later that in this picture, our corresponding definitions of classical and quantum correlations coincide precisely with the usual ones [2].

So, generalizing and reinterpreting some of Bell's work is what we set out to do in this paper. The questions that naturally arise are of this kind: under which conditions can a given correlation be explained in terms of hidden variables on a given causal structure? We are not able to answer this question in this paper and we suspect that it is in general very difficult, but we provide a set of basic definitions and results for considering it.

3. Correlations

So what are these "correlations" between observations which we were talking about? As mentioned before, we imagine that at every node $v \in V$, there is an outcome $o[v]$ describing the result of the corresponding observation and taking values in a finite set of possible values denoted O_v. It takes a concrete value in each run of the experiment, but possibly not the same value in each run, so that it needs to be described in terms of a probability distribution over outcomes. We also write $o[V]$ for the vector or tuple $(o[v])_{v\in V}$ consisting of all the outcomes $o[v]$ for all nodes $v \in V$. In each run of the experiment, we obtain one concrete tuple of values for $o[V]$, so that we obtain a probability distribution $P(o[V])$, which is a joint distribution of all the individual outcome variables $o[v]$. If $U \subseteq V$ is any subset of the nodes, then we similarly use notation like $o[U]$ and $P(o[U])$ for the corresponding tuples and their observed distribution.

Now, under which condition is a given joint distribution $P(o[V])$ explainable in terms of hidden variables on a given causal structure? This is a very difficult question in general and we are very far from having an answer, but there is an obvious necessary condition: if two subsets $V_1, V_2 \subseteq V$ have distinct causal past, then $o[V_1]$ and $o[V_2]$ should be independent. More generally:

3.1. Definition. *$P(o[V])$ is a **correlation** if for any collection of subsets $V_1, \dots, V_m \subseteq V$ with disjoint causal past, i.e. $\mathrm{pst}(V_i) \cap \mathrm{pst}(V_j) = \emptyset$ for $i \neq j$, the distribution factors as*

$$P(o[V_1] \dots o[V_m]) = P(o[V_1]) \cdots P(o[V_m]). \tag{3.1}$$

This is the basic necessary condition that we will use. From now on, we restrict our considerations to those joint distributions $P(o[V])$ that are correlations. No other joint distributions of outcomes will be considered on a given causal structure. When we define classical correlations, it will be our first result (Proposition 4.2) to show that every classical correlation is actually a correlation, and similarly in the quantum case.

In deriving the concrete form of the constraints (3.1) for a given (candidate) causal structure, it is sufficient to postulate (3.1) only for $m = 2$, i.e. when one only considers two sets U and W with disjoint causal past, in which case the condition is

$$P(o[U]o[W]) = P(o[U])\, P(o[W]). \tag{3.2}$$

The reason for the equivalence is that one can derive (3.1) from this assumption recursively. For example, for three sets V_1, V_2, and V_3 with disjoint causal past, one gets

$$\begin{aligned} P(o[V_1]o[V_2]o[V_3]) &= P(o[V_1]o[V_2 \cup V_3]) \overset{(3.2)}{=} P(o[V_1])\, P(o[V_2 \cup V_3]) \\ &= P(o[V_1])\, P(o[V_2]o[V_3]) \overset{(3.2)}{=} P(o[V_1])\, P(o[V_2])\, P(o[V_3]). \end{aligned}$$

Here, the second equation uses the fact that V_1 and $V_2 \cup V_3$ have disjoint causal past, which follows from the assumption that all the V_i have pairwise disjoint causal past.

Moreover, in (3.2) it is enough to consider those U and W which are **maximal** in the sense that there is no node outside of these sets which could be added to either of them, while still preserving the condition of disjoint past. The reason for this equivalence is that for any pair U, W, there is a maximal one U', W' with $U \subseteq U'$ and $W \subseteq W'$. Now the assumption

$$P(o[U']o[W']) = P(o[U'])\, P(o[W'])$$

implies the desired equation (3.3) upon taking the marginal of both sides over all those variables which are contained in U' or W', but not in U or W. Note that maximality requires in particular

that $\mathrm{pst}(U) = U$ and $\mathrm{pst}(W) = W$, since otherwise U or W could be enlarged. Thereby, we have shown:

3.2. Proposition. *$P(o[V])$ is a correlation if and only if for any maximal pair of subsets $U, W \subseteq V$ with $\mathrm{pst}(U) \cap \mathrm{pst}(W) = \emptyset$, one has*

$$P(o[U]o[W]) = P(o[U])\, P(o[W]). \tag{3.3}$$

For any given causal structure, it is instructive to write down a minimal set of conditions on $P(o[V])$ which guarantee that it is a correlation. We now carry this out in the familiar case of a Bell scenario. Instead of considering only the 2-party scenario of Figure 1, we generalize it to the ***n*-party Bell scenario** in which the two "arms" $\{a, x\}$ and $\{b, y\}$ are replaced by an arbitrary number of arms $\{a_1, x_1\}, \ldots, \{a_n, x_n\}$. More precisely, the n-party Bell scenario has $2n + 1$ nodes: a "source" s, "settings" $x_1, \ldots, x_n$ and "outcomes" $a_1, \ldots, a_n$. For each party i we have $s \to a_i$ and $x_i \to a_i$, and there are no other edges. Our first result about the n-party Bell scenario is that correlations on this causal structure can be characterized like this:

3.3. Proposition. *In the n-party Bell scenario, P is a correlation if and only if it satisfies the* ***free will equation***

$$P(o[x_1 \ldots x_n s]) = P(o[x_1]) \cdots P(o[x_n])\, P(o[s]) \tag{3.4}$$

and the ***no-signalling equation***

$$P(o[a_1 \ldots \cancel{a_i} \ldots a_n x_1 \ldots x_n s]) = P(o[x_i]) P(o[a_1 \ldots \cancel{a_i} \ldots a_n x_1 \ldots \cancel{x_i} \ldots x_n s]) \tag{3.5}$$

for every party $i = 1, \ldots, n$.

If one makes use of (3.4), then the no-signalling equation (3.5) can be seen to be equivalent to the more conventional form

$$P(o[a_1 \ldots \cancel{a_i} \ldots a_n] \,|\, o[x_1 \ldots x_n s]) = P(o[a_1 \ldots \cancel{a_i} \ldots a_n] \,|\, o[x_1 \ldots \cancel{x_i} \ldots x_n s]),$$

which is usually interpreted as stating that the conditional probability on the left is independent of the particular value of $o[x_i]$. Also, it is usually assumed in the Bell scenario that the source s does not have an outcome; in our formalism, this corresponds to taking the source's outcome $o[s]$ to be always the same. In an actual experiment, however, the source typically does have an outcome, corresponding to whether it successfully produced the required pair of entangled particles or not. In this sense, our approach automatically treats a Bell scenario in a more realistic manner than is usually done. Another difference is that a correlation in our approach does include chosen probabilities $P(o[x_i])$ for the "settings" $o[x_i]$. Again, this seems more realistic to us: in an actual experiment, the $o[x_i]$ have particular distributions, and these are even relevant in practical information processing applications like device-independent randomness expansion, where strong limitations on the $P(o[x_i])$ need to be imposed [6]. Another example is device-independent quantum key distribution [4], where protocol performance also depends on the $P(o[x_i])$, even in such a way that highly non-uniform distributions are optimal.

Proof. We begin with the "only if" part. To establish (3.4) from the assumption that $P(o[V])$ is a correlation, use

$$V_1 = \{x_1\}, \quad \ldots, \quad V_n = \{x_n\}, \quad V_{n+1} = \{s\}$$

in (3.1). To establish (3.5), put similarly

$$V_1 = \{x_i\}, \qquad V_2 = \{a_1 \ldots \cancel{a_i} \ldots a_n x_1 \ldots \cancel{x_i} \ldots x_n s\}.$$

In both cases, it is straightforward to check that the condition of having disjoint causal past holds.

For the "if" part, suppose (3.4) and (3.5) hold. Thanks to Proposition 3.2, it suffices to check the correlation condition for two maximal disjoint-past sets U and W. By $\mathrm{pst}(U) \cap \mathrm{pst}(W) = \emptyset$, at least one of $\mathrm{pst}(U)$ and $\mathrm{pst}(W)$ does not contain s; without loss of generality, we take this to be U. Therefore, by maximality, s must be contained in W. Then since U and W have disjoint causal past, U cannot contain s or any a_i, so $U = \{\, x_i \mid i \in I \,\}$ for some $I \subseteq \{1, \ldots, n\}$. Again by maximality, we conclude that $W = \{\, a_i, x_i \mid i \notin I \,\} \cup \{\, s \,\}$. If we write $\bar{I} := \{1, \ldots, n\} \backslash I$, then the equation that we have to derive from (3.4) and (3.5) is that

$$P(o[a_{\bar{I}} x_I x_{\bar{I}} s]) = P(o[x_I])\, P(o[a_{\bar{I}} x_{\bar{I}} s]),$$

for an arbitrary set of parties $I \subseteq \{1, \ldots, n\}$. This equation can be interpreted as the analogue of (3.5) in which the individual party i is replaced by the set I. That this indeed follows from the assumptions is analogous to a well-known fact in the conventional formalism of no-signaling boxes, where no-signaling across a bipartition I vs $\bar{I}$ follows from no-signaling with respect to each party individually [**25**]. With our definitions, this derivation can be carried out by induction on the size of I: if I contains just one party, then our target equation coincides with the assumption (3.5). For the induction step, we assume that the equation holds for I and prove that it also holds for $I \cup \{\, j \,\}$ in place of I, where j is an arbitrary party not in I. This works as follows:

$$\begin{aligned} P(o[a_{\bar{I}\backslash\{j\}} x_{I\cup\{j\}} x_{\bar{I}\backslash\{j\}} s]) &= P(o[a_{\bar{I}\backslash\{j\}} x_I x_{\bar{I}} s]) \\ &= P(o[x_I])\, P(o[a_{\bar{I}\backslash\{j\}} x_{\bar{I}} s]) \\ &\overset{(3.5)}{=} P(o[x_I])\, P(o[x_j])\, P(o[a_{\bar{I}\backslash\{j\}} x_{\bar{I}\backslash\{j\}} s]) \\ &\overset{(3.4)}{=} P(o[x_{I\cup\{j\}}])\, P(o[a_{\bar{I}\backslash\{j\}} x_{\bar{I}\backslash\{j\}} s]), \end{aligned}$$

as was to be shown. Here, the second step uses the induction assumption, marginalized over $o[a_j]$. The third and fourth steps use similar marginalizations of (3.5) and (3.4). □

Finally, it is worth noting that the set of all correlations on a given graph G does not depend on all details of the graph itself, but only on the partial order "$\rightsquigarrow$" induced by G, as discussed in Section 2. The reason is that in order to apply Definition 3.1, all that one needs to know about the causal structure is which events lie in the causal past of which other ones.

4. Classical correlations

As discussed in the introduction and the previous section, we want to assume that every node $v \in V$ carries a random variable $o[v]$, which represents the outcome of some probing of spacetime conducted at v. The basic assumption that we make is that these probings may not constitute a complete and faithful representation of the physical systems propagating through these events, and in addition the observations may modify those systems in an unknown and arbitrary manner. More concretely, it could be the case that there are additional "hidden variables" necessary to describe the systems under consideration, and the observations only reveal partial information about the values of these hidden variables. We believe that this is a hypothesis which naturally arises in many situations in the sciences: for example, we can measure action potentials in a biological neural network and try to infer the causal structure from the correlations behind these action potentials, but how would we know that there is no other form of information being transmitted between the neurons? If we measure a certain number of meteorological parameters in the atmosphere at different points in space and time, how can we be sure that we are not missing relevant information like the concentration of a certain aerosol, which may have a crucial effect on cloud formation? For these reasons, we allow for the propagation of hidden information and for the possibility that the observations modify this hidden information in an unknown manner.

There are two ways to make this idea precise:

- The hidden variables live on the *edges* of the graph.

  This is the approach that we pursue in this section. The local structure of a node v is illustrated in Figure 6: the ingoing edges $e \in E$ carry hidden variables λ_e, which determine the outcome $o[v]$ and the outgoing hidden variables via a certain conditional distribution $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$[1]. In terms of a computational interpretation, the node v operates like an **information processing gate** taking the incoming hidden variables $\lambda_{\mathrm{in}(v)}$ and turning them into outgoing hidden variables $\lambda_{\mathrm{out}(v)}$ together with an outcome $o[v]$. In terms of a physical interpretation, we think that it is appropriate to have the hidden variables live on the edges: after all, a hidden variable is supposed to represent the state of a physical system, and since we interpret our causal structures in terms of spacetime pictures, these physical systems should be described by their one-dimensional worldlines. Since we allow the systems to be modified in an arbitrary way by the observations, the conditional distribution $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$ may be completely arbitrary. In particular, the outgoing systems may be of a different kind than the incoming ones, and we do not make any sort of "noninvasiveness" assumption [26], and hence our formalism does not comprise Leggett-Garg scenarios.
- The hidden variables live on the *nodes* of the graph.

  We develop this approach in Section 7 and show it to be equivalent to the previous one. While it is less susceptible to a physical or information processing interpretation, it is very natural from the mathematical point of view and fits in nicely with the standard approach to causal inference in terms of Bayesian networks [13]. It captures the idea of assuming that there may be hidden information "behind" each node.

We now develop the first approach in some detail.

4.1. Definition. *$P(o[V])$ is a **classical correlation** if for every $e \in E$ there is a hidden variable λ_e given in terms of a probability space (X_e, Σ_e), and for every $v \in V$ there is a conditional*

[1]Here, we also write $\lambda_{\mathrm{out}(v)}$ as shorthand for the tuple of variables $(\lambda_e)_{e\in\mathrm{out}(v)}$ etc.

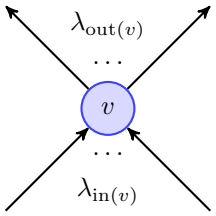

FIGURE 6. Intended local structure of a node v.

distribution

$$P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \tag{4.1}$$

such that

$$P(o[V]) = \int_{\lambda_E} \prod_{v\in V} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}). \tag{4.2}$$

There are many possible variations on this definition having to do with the precise mathematical definitions used in the case that the hidden variables may take infinitely many values, which we explicitly allow. The definitions that we work with are explained in Appendix A; we will not consider other possible definitions in this work. We suspect that the set of classical correlations does not depend on which definitions are used, but we do not have a proof of this. The readers not interested in non-discrete hidden variables may want to restrict themselves to taking the X_e to be finite or countable sets.

Some further explanation of (4.2) is in order. For any concrete value of the tuple of outcomes $o[V]$, the right-hand side evaluates to a non-negative number as follows. Any $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$ is a $\lambda_{\mathrm{in}(v)}$-dependent subnormalized measure on the product space $X_{\mathrm{out}(v)} = \prod_{e\in\mathrm{out}(v)} X_e$. The product in (4.2) hence yields a measure on the total product space $X_E = \prod_{e\in E} X_e$. Integrating this measure results in a number, corresponding to its normalization, which is the joint probability of getting the joint outcome $o[V]$. Strictly speaking, making sense of this product of conditional distributions and the integral requires the methods of Section 5, since it is not obvious how to show that the integrand does indeed define a measure. We ignore this technical point for the time being.

As the terminology suggests, a classical correlation is indeed a correlation in the sense of Section 3.

4.2. PROPOSITION. *If $P(o[V])$ is a classical correlation, then it is a correlation.*

PROOF. It is already non-trivial to verify that $P(o[V])$ is actually a probability distribution, i.e. satisfies the normalization equation $\sum_{o[V]} P(o[V]) = 1$.

We start by showing that if $U \subseteq V$ is some set of nodes with $\mathrm{pst}(U) = U$, then the marginal $P(o[U])$ of (4.2) can be computed like this:

$$P(o[U]) = \int_{\lambda_{\mathrm{out}(U)}} \prod_{v\in U} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}). \tag{4.3}$$

Here, $\mathrm{out}(U) = \bigcup_{v\in U} \mathrm{out}(v)$, so that the integration variables $\lambda_{\mathrm{out}(U)}$ comprise all those λ_e with $\mathrm{src}(e) \in U$. This formula should make intuitive sense: in order to determine the distribution of outcomes in $\mathrm{pst}(U) = U$, it should be sufficient to consider only the nodes and edges in U itself, since no other events or hidden variables can influence the variables in U.

We prove (4.3) by induction on the size of $V\backslash U$. The base case is $U = V$, for which the formula is precisely the definition (4.2). For the induction step, we need to prove that the formula holds for a given U with $\mathrm{pst}(U) = U$, when assuming that it holds for all bigger U' in place of U which likewise satisfy $\mathrm{pst}(U') = U'$. For the given U, we can choose a node w with $\mathrm{pa}(w) \subseteq U$; upon putting $U' := U \cup \{w\}$, it is then guaranteed that $\mathrm{pst}(U') = U'$ as well. We therefore obtain, by the induction assumption,

$$\begin{aligned}
P(o[U]) &= \sum_{o[w]} P(o[U]o[w]) = \sum_{o[w]} \int_{\lambda_{\mathrm{out}(U)\cup\mathrm{out}(w)}} P(o[w]\lambda_{\mathrm{out}(w)}|\lambda_{\mathrm{in}(w)}) \prod_{v\in U} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \\
&= \int_{\lambda_{\mathrm{out}(U)}} \int_{\lambda_{\mathrm{out}(w)}} \sum_{o[w]} P(o[w]\lambda_{\mathrm{out}(w)}|\lambda_{\mathrm{in}(w)}) \prod_{v\in U} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \\
&= \int_{\lambda_{\mathrm{out}(U)}} \left(\prod_{v\in U} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \right) \int_{\mathrm{out}(w)} \sum_{o[w]} P(o[w]\lambda_{\mathrm{out}(w)}|\lambda_{\mathrm{in}(w)}) \\
&= \int_{\lambda_{\mathrm{out}(U)}} \prod_{v\in U} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}).
\end{aligned}$$

Here, in the last step we have used the normalization of probability for the conditional distribution $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$, which states precisely that the integral-sum on the right evaluates to 1.

This completes the proof of (4.3). In the special case $U = \emptyset$, the right-hand side of (4.3) is an empty product and therefore evaluates to 1, which proves that $P(o[V])$ is indeed a probability distribution.

To keep the notation simple, we prove the "binary" correlation equation (3.3) with the assumption $\mathrm{pst}(U) = U$ and $\mathrm{pst}(W) = W$, which we already know to suffice. Then we have $\mathrm{pst}(U \cup W) = \mathrm{pst}(U) \cup \mathrm{pst}(W) = U \cup W$, and therefore we obtain by (4.3),

$$\begin{aligned}
P(o[U]o[W]) &= \int_{\lambda_{\mathrm{out}(U\cup W)}} \prod_{v\in U\cup W} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \\
&= \int_{\lambda_{\mathrm{out}(U)}} \int_{\lambda_{\mathrm{out}(W)}} \left(\prod_{v\in U} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \right) \cdot \left(\prod_{v\in W} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \right) \\
&= \left(\int_{\lambda_{\mathrm{out}(U)}} \prod_{v\in U} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \right) \cdot \left(\int_{\lambda_{\mathrm{out}(W)}} \prod_{v\in W} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \right) \\
&= P(o[U])\, P(o[W]),
\end{aligned}$$

as was to be shown. □

In the case of Bell scenarios, classical correlations correspond exactly to those conditional probability distributions that are usually called "local" [2]:

4.3. Proposition. *In the n-party Bell scenario, P is a classical correlation if and only if it satisfies the free will equation (3.4), and the corresponding conditional probabilities are of the form*

$$P(o[a_1 \dots a_n]|o[x_1 \dots x_n s]) = \int_\lambda P(o[a_1]|o[x_1]\lambda) \cdots P(o[a_n]|o[x_n]\lambda) P(\lambda|o[s]) \tag{4.4}$$

for some single hidden variable λ living on some measurable space (X, Σ).

Here, we implicitly also assume that only those values of $o[x_i]$ and $o[s]$ are considered for which $P(o[x_i]) > 0$ and $P(o[s]) > 0$, since otherwise conditioning on these variables would not make sense.

PROOF. We begin with the "only if" part. To establish (4.4) for a given classical correlation, we define the conditional distributions occurring on the right-hand side of (4.4) by identifying $\lambda := \lambda_{\mathrm{out}(s)}$ and putting

$$P(\lambda|o[s]) := \frac{P(o[s]\lambda_{\mathrm{out}(s)})}{P(o[s])},$$

$$P(o[a_i]|o[x_i]\lambda_{\mathrm{out}(s)}) := \int_{\lambda_{x_i\to a_i}} \frac{P(o[a_i]|\lambda_{s\to a_i}\lambda_{x_i\to a_i})P(o[x_i]\lambda_{x_i\to a_i})}{P(o[x_i])}$$

where the (conditional) distributions in the numerators of the right-hand sides are all part of the given data which realizes P as a classical correlation, while those in the denominators are marginals thereof. Using the free will equation (3.4), we have

$$P(o[a_1\dots a_n]|o[x_1\dots x_n s]) = \frac{P(o[a_1\dots a_n x_1\dots x_n s])}{P(o[x_1])\cdots P(o[x_n])P(o[s])} \tag{4.5}$$

which, when the numerator is expanded using (4.2), gives (4.4).

Now for the "if" part. Suppose (3.4) and (4.4) hold. We establish that this is a classical correlation as follows. Let each $\lambda_{s\to a_i}$ be a copy of λ. Let $P(s\lambda_{\mathrm{out}(s)}) := P(\lambda|s)\,P(s)$. We take $\lambda_{x_i\to a_i}$ to be $o[x_i]$ itself, meaning that $P(o[x_i]\lambda_{x_i\to a_i}) := P(o[x_i])\cdot\delta_{o[x_i],\lambda_{x_i\to a_i}}$. Finally, let $P(o[a_i]|\lambda_{x_i\to a_i}\lambda_{s\to a_i}) := P(o[a_i]|o[x_i]\lambda)$. Multiplying (4.4) by (3.4) then gives (4.2). □

A similar argument will work in the quantum case; see Proposition 6.5.

Bell scenarios have the desirable feature that the resulting conditional distributions of the form (4.4) form a convex polytope known as the **local polytope** [2]. Checking whether a given correlation is classical then simply boils down to asking whether the associated conditional distribution lies in the local polytope. However, for general causal structures, determining whether a given correlation is classical is a very difficult problem. A good example is the triangle scenario displayed in Figure 4, which has been studied in [10, 16, 17] for the special case that the outcomes $o[x]$, $o[y]$ and $o[z]$ on the first layer of nodes are deterministic variables. In this situation, the only known generally applicable conditions for a correlation to be classical are certain entropic inequalities; see [10, Lemma 2.14] and [17, Sec. VII.B]. One of the problematic aspects is that the set of classical correlations is not convex. In fact, this already holds in the case of a Bell scenario: although the set of classical correlations is a convex polytope on the level of the conditional distributions (4.4), convexity fails as soon as one works in the space of "unconditional" distributions $P(o[a_1\dots a_n x_1\dots x_n s])$.

So it is an important problem to understand the structure of classical correlations better. One relevant question is whether one can bound the maximally necessary sizes of the hidden variable spaces (X_e, Σ_e) in terms of the number of outcomes of each $o[v]$. In Bell scenarios, this was shown to be possible by Fine [27], and it again has to do with the fact that the associated conditional distributions form a convex polytope. But in general, it is not even clear whether one must allow the hidden variable spaces (X_e, Σ_e) to be continuous or even *bigger than continuous* in order to obtain all classical correlations. What we have been able to prove is that finite hidden variable spaces are sufficient at least in an approximate sense:

4.4. Theorem. *If $P(o[V])$ is a classical correlation on an arbitrary causal structure G and $\varepsilon > 0$, then there is another classical correlation $P'(o[V])$ on G with $|P(o[V]) - P'(o[V])| < \varepsilon$ for any value of $o[V]$ and such that the probability spaces (X'_e, Σ'_e) occurring in the hidden variable model for P' are all finite.*

It remains open whether this also holds without the approximation, i.e. whether any classical correlation can be exactly represented in terms of finite hidden variable spaces. We suspect that this is not the case, but we do not have any examples. If one generalizes our definitions to the case of countably infinite G, as we do informally in the discussion opening Section 7, then there are well-known examples [**28**, Sec. 3.6].

Proof. We construct, for every $\delta > 0$, a classical correlation P'_δ realized by hidden variables with only finitely many values such that $|P(o[V]) - P'_\delta(o[V])|$ converges to 0 as $\delta \to 0$. We perform this construction by induction on $|V|$, the number of nodes of the given graph.

For $|V| = 1$, we have no edges and therefore no hidden variables. Hence we can simply put $P'_\delta(o[V]) := P(o[V])$ for any δ, and there is nothing to prove.

For the induction step, we start with a graph $G = (V, E)$ and a classical correlation $P(o[V])$ realized by hidden variables λ_e as in the classical correlation equation (4.2), where the hidden variable spaces (X_e, Σ_e) may be of arbitrary size. We then choose a node $w \in V$ with $\mathrm{out}(w) = \emptyset$; such a node exists by the assumptions of finiteness and acyclicity. Taking out the node w together with all the edges with target w gives a smaller graph to which the induction hypothesis can be applied, but we need to be careful to do this in the right way. Before doing so, we need a bit more preparation.

We first consider the conditional distribution $P(o[w]|\lambda_{\mathrm{in}(w)})$ at the chosen node w. We choose a finite subset of $[0, 1]$ such that any number in $[0, 1]$ has distance at most δ from this finite subset. Then we can replace the given conditional distribution $P(o[w]|\lambda_{\mathrm{in}(w)})$ by a perturbation $\hat{P}_\delta(o[w]|\lambda_{\mathrm{in}(w)})$ which only takes values in this finite subset and such that

$$|\hat{P}_\delta(o[w]|\lambda_{\mathrm{in}(w)}) - P(o[w]|\lambda_{\mathrm{in}(w)})| < \delta \qquad \forall\, o[w], \lambda_{\mathrm{in}(w)}. \tag{4.6}$$

Then, for any given tuple of values for $\lambda_{\mathrm{in}(w)}$, there are only finitely many possibilities for the conditional distribution $\hat{P}_\delta(o[w]|\lambda_{\mathrm{in}(w)})$. Hence we can consider the family of these distributions, as indexed by $\lambda_{\mathrm{in}(w)}$, as a measurable function from $X_{\mathrm{in}(w)} = \prod_{e \in \mathrm{in}(w)} X_e$ to some finite set S. We consider the product measurable space $X_{\mathrm{in}(w)}$ to be equipped with the corresponding marginal of the total measure

$$P(o[V]\lambda_E) = \prod_{v \in V} P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}), \tag{4.7}$$

which determines the distribution of the $\lambda_{\mathrm{in}(w)}$. By virtue of Proposition B.2, we can replace the measurable spaces (X_e, Σ_e) for all $e \in \mathrm{in}(w)$ by finite sets X'_e carrying hidden variables μ_e. More concretely, the X'_e are coarse-grainings of the X_e arising via functions $f_e : X_e \to X'_e$ and come themselves equipped with a function

$$P'_\delta \,:\, X'_{\mathrm{in}(w)} \longrightarrow S$$

which assigns to every tuple $\mu_{\mathrm{in}(w)} \in X'_{\mathrm{in}(w)}$ a distribution $P'_\delta(o[w]|\mu_{\mathrm{in}(w)})$ having the property that

$$P'_\delta\big(o[w]|f_{\mathrm{out}(w)}(\lambda_{\mathrm{out}(w)})\big) = \hat{P}_\delta(o[w]|\lambda_{\mathrm{in}(w)}) \tag{4.8}$$

holds with probability greater than $1 - \delta$ with respect to the original distribution $P(\lambda_{\mathrm{in}(w)})$ as the marginal of (4.7).

In order to apply the induction hypothesis, we need to modify the correlation a bit by regarding each new hidden variable μ_e for $e \in \mathrm{in}(w)$ as part of the outcome $o[\mathrm{src}(e)]$ associated to the source node $\mathrm{src}(e)$. In other words, we define new outcomes:

$$\tilde{O}_v := \begin{cases} O_v \times X'_{v\to w} & \text{if } v \to w, \\ O_v & \text{if } v \not\to w, \end{cases} \qquad \tilde{o}[v] := \begin{cases} (o[v], \mu_{v\to w}) & \text{if } v \to w, \\ o[v] & \text{if } v \not\to w. \end{cases}$$

Now there is an obvious classical correlation on the remainder graph $G\backslash\{w\}$ corresponding to a distribution of the new outcomes $\tilde{o}[V]$: at those v with $v \not\to w$, we retain the previous conditional distribution $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$; at those v with $v \to w$, we compose the old $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$ with the f_e from above, which gives a new conditional distribution $P(o[v]\mu_e\lambda_{\mathrm{out}(v)\backslash\{w\}}|\lambda_{\mathrm{in}(v)})$, in which the first two components are now considered to be forming the new outcome $\tilde{o}[v]$. This data defines a new classical correlation $P(\tilde{o}[V])$ on the smaller graph $G\backslash\{w\}$.

We now apply the induction hypothesis to *this* classical correlation. This way, we obtain finite sets X'_e carrying hidden variables μ_e for all edges $e \in E$ with $\mathrm{tar}(e) \neq w$, where the new variables μ_e have a priori nothing to do with the previous $\mu_{\mathrm{in}(w)}$ except for sharing the same letter, and conditional probabilities $P'_\delta(\tilde{o}[v]\mu_{\mathrm{out}(v)}|\mu_{\mathrm{in}(v)})$ such that the difference

$$|P'_\delta(o[V\backslash\{w\}]\,\mu_{\mathrm{in}(w)}) - P(o[V\backslash\{w\}]\,\mu_{\mathrm{in}(w)})| = |P'_\delta(\tilde{o}[V\backslash\{w\}]) - P(\tilde{o}[V\backslash\{w\}])| \tag{4.9}$$

tends to 0 as $\delta \to 0$.

Taking all this together, we will obtain the desired bounds on the difference

$$\begin{aligned} &P(o[V]) - P'_\delta(o[V]) \\ &\quad = P(o[V]) - \sum_{\mu_{\mathrm{in}(w)}} P'_\delta(o[w]|\mu_{\mathrm{in}(w)})P(o[V\backslash\{w\}]\mu_{\mathrm{in}(w)}) \\ &\quad + \sum_{\mu_{\mathrm{in}(w)}} \left(P'_\delta(o[w]|\mu_{\mathrm{in}(w)})P(o[V\backslash\{w\}]\mu_{\mathrm{in}(w)}) - P'_\delta(o[w]|\mu_{\mathrm{in}(w)})P'_\delta(o[V\backslash\{w\}]\mu_{\mathrm{in}(w)})\right), \end{aligned}$$

where we have inserted the first sum such that it cancels with the first part of the second sum. The second sum can be bounded in terms of (4.9): by making the δ used in the application of the induction hypothesis arbitrarily small, we can also make the second sum arbitrarily small. Thus it remains to bound the difference between $P(o[V])$ and the first sum, which can be done like this:

$$\begin{aligned} &\left|P(o[V]) - \sum_{\mu_{\mathrm{in}(w)}} P'_\delta(o[w]|\mu_{\mathrm{in}(w)})P(o[V\backslash\{w\}]\mu_{\mathrm{in}(w)})\right| \\ &\qquad = \left|P(o[V]) - \int_{\lambda_{\mathrm{in}(w)}} P'_\delta(o[w]|f_{\mathrm{in}(w)}(\lambda_{\mathrm{in}(w)}))P(o[V\backslash\{w\}]\lambda_{\mathrm{in}(w)})\right| \\ &\qquad \overset{(4.8)}{\leq} \left|P(o[V]) - \int_{\lambda_{\mathrm{in}(w)}} \hat{P}_\delta(o[w]|\lambda_{\mathrm{in}(w)})P(o[V\backslash\{w\}]\lambda_{\mathrm{in}(w)})\right| + \delta \\ &\qquad \overset{(4.6)}{\leq} \left|P(o[V]) - \int_{\lambda_{\mathrm{in}(w)}} P(o[w]|\lambda_{\mathrm{in}(w)})P(o[V\backslash\{w\}]\lambda_{\mathrm{in}(w)})\right| + 2\delta \\ &\qquad = 2\delta. \end{aligned}$$

□

In Bell scenarios, it is well-known that the "information processing" at the measurement nodes a_i may be assumed to be deterministic without loss of generality [27], in the sense that any classical correlation has a hidden variable model satisfying this property. For general causal structures, it also seems very plausible that it should be possible to assume that the information processing at each node v with $\mathrm{pa}(v) \neq \emptyset$ is deterministic. The idea here is that if v has at least one parent node, then the randomness needed can be "pushed back" in the sense of generating the randomness at the parent node and then transporting it from there to v by augmenting the hidden variable on that edge with this additional information. If one repeats this process of pushing randomness generation back to the parents nodes, then one eventually ends up with a classical model in which all information processing is deterministic, except at those nodes that do not have any parents, which is where all the randomness originates.

More precisely, the hypothesis is this:

4.5. Conjecture. *If $P(o[V])$ is a classical correlation, then $P(o[V])$ can be represented in the form (4.2) in such a way that all conditional probabilities are deterministic at all nodes v with $\mathrm{in}(v) \neq \emptyset$, meaning that $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$ is a Dirac measure on $O_v \times X_{\mathrm{out}(v)}$ for all values of $\lambda_{\mathrm{in}(v)}$.*

Due to the measure-theoretical technicalities involved, we only have a partial proof of this conjecture covering the case in which all hidden variables take only finitely many values. In this case, the statement can be reformulated as saying that $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}) \in \{0, 1\}$ for any values of the variables involved. Our proof is based on exactly the intuitive idea of repeatedly pushing the randomness back to parent nodes:

Partial Proof. As in the proof of Theorem 4.4, we use induction on the size of G. When G has only vertex, then there is no hidden variable and hence nothing to prove.

For the induction step, we start with a graph $G = (V, E)$ and a classical correlation $P(o[V])$ on G realized by hidden variables λ_e living on finite sets X_e and conditional distributions $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$ which are, due to the finiteness assumption, simply stochastic matrices. If there is no node with a parent, then the graph has no edges at all, and again there is nothing to prove. Otherwise, we pick a vertex w with $\mathrm{pa}(w) \neq \emptyset$, but $\mathrm{ch}(w) = \emptyset$; since G is finite and acyclic, there is at least one such vertex which can be found e.g. by starting at any node with an incoming edge and following outgoing edges from node to node until one reaches a node without any outgoing edges. From now on, we keep such a node w fixed.

The induction step consists in "pushing back" the randomness inherent in the conditional distribution $P(o[w]|\lambda_{\mathrm{in}(w)})$ to a parent node of w, and applying the induction hypothesis afterwards. While this should be intuitively plausible, we will see that the technical details are quite demanding.

So to perform this induction step, pick an arbitrary parent $u \in \mathrm{pa}(w)$, which is where the randomness will be pushed back to. Doing so means that the hidden variable space $X_{u\to w}$ will have to be enlarged, and the appropriate definition turns out to be

$$X'_{u\to w} := O_w^{X_{\mathrm{in}(w)}} \times X_{u\to w},$$

where the first factor is the set of all functions $X_{\mathrm{in}(w)} \to O_w$. The given conditional distribution $P(o[w]|\lambda_{\mathrm{in}(w)})$ can be regarded as a distribution on this function space by assigning to a function $f : X_{\mathrm{in}(w)} \to O_w$ the probability

$$\hat{P}(f) := \prod_{\lambda_{\mathrm{in}(w)} \in X_{\mathrm{in}(w)}} P(f(\lambda_{\mathrm{in}(w)})|\lambda_{\mathrm{in}(w)}).$$

Intuitively, a function $f : X_{\mathrm{in}(w)} \to O_w$ can also be regarded as a list of elements of O_w indexed by $X_{\mathrm{in}(w)}$, and then this definition says that all the elements of this list are selected independently of each other according to their respective distributions determined by the given $P(o[w]|\lambda_{\mathrm{in}(w)})$.

If we denote the new hidden variable with the new domain $X'_{u\to w}$ by $\lambda'_{u\to w}$, then we can define a new conditional distribution $P'(o[w]|\lambda'_{u\to w}\lambda_{\mathrm{in}(w)\setminus\{u\to w\}})$ by writing $\lambda'_{u\to w}$ as a pair $\lambda'_{u\to w} = (f, \lambda_{u\to w})$ and putting

$$P'(o[w]|(f,\lambda_{u\to w})\lambda_{\mathrm{in}(w)\setminus\{u\to w\}}) := \begin{cases} 1 & \text{if } f(\lambda_{\mathrm{in}(w)}) = o[w], \\ 0 & \text{otherwise,} \end{cases}$$

where $\lambda_{\mathrm{in}(w)}$ stands for the tuple of hidden variables consisting of the $\lambda_{\mathrm{in}(w)\setminus\{u\to w\}}$ together with the additional $\lambda_{u\to w}$ which is part of $\lambda'_{u\to w}$. By construction, P' is deterministic information processing: the node w receives a collection of incoming hidden variables $\lambda_{\mathrm{in}(w)}$ as well as a function $f : X_{\mathrm{in}(w)} \to O_w$ and does nothing but to apply the latter to the former in a deterministic manner.

In order to define a classical hidden variable model using the new hidden variable space $X'_{u\to w}$, we also need to modify the conditional distribution $P(o[u]\lambda_{\mathrm{out}(u)}|\lambda_{\mathrm{in}(u)})$ to some $P'(o[u]\lambda'_{u\to w}\lambda_{\mathrm{out}(u)\setminus\{u\to w\}}|\lambda_{\mathrm{in}(u)})$. Since $\lambda'_{u\to w} = (f, \lambda_{u\to w})$, we can do this by simply putting

$$P'(o[u]f\lambda_{\mathrm{out}(u)}|\lambda_{\mathrm{in}(u)}) := \hat{P}(f)\cdot P(o[u]\lambda_{\mathrm{out}(u)}|\lambda_{\mathrm{in}(u)}),$$

where $\hat{P}(f)$ is as above. In other words, the function f is chosen already at u, randomly and independently of any of the other hidden variables. In this sense, the randomness has been "pushed back" from the node w to the node u.

In order to show that the two new conditional distributions P' at nodes u and w indeed give rise to the same overall classical correlation, it is enough to show that the resulting conditional distribution $P'(o[uw]\lambda_{\mathrm{out}(u)\setminus\{u\to w\}}|\lambda_{\mathrm{in}(u)\cup\mathrm{in}(w)\setminus\{u\to w\}})$ coincides with the original one, since this is the only part of the graph to which we have applied any modifications. Indeed,

$$\begin{aligned}
&P'(o[uw]\lambda_{\mathrm{out}(u)\setminus\{u\to w\}}|\lambda_{\mathrm{in}(u)\cup\mathrm{in}(w)\setminus\{u\to w\}}) \\
&\quad= \sum_{f,\lambda_{u\to w}} P'(o[w]|f\lambda_{\mathrm{in}(w)})\, P'(o[u]f\lambda_{\mathrm{out}(u)}|\lambda_{\mathrm{in}(u)}) \\
&\quad= \sum_{f,\lambda_{u\to w}} \delta_{f(\lambda_{\mathrm{in}(w)}),o[w]}\cdot\hat{P}(f)\cdot P(o[u]\lambda_{\mathrm{out}(u)}|\lambda_{\mathrm{in}(u)}) \\
&\quad= \sum_{\lambda_{u\to w}} P(o[u]\lambda_{\mathrm{out}(u)}|\lambda_{\mathrm{in}(u)}) \quad \sum_{f}\delta_{f(\lambda_{\mathrm{in}(w)}),o[w]}\prod_{\mu_{\mathrm{in}(w)}} P(f(\mu_{\mathrm{in}(w)})|\mu_{\mathrm{in}(w)}) \\
&\quad\overset{(*)}{=} \sum_{\lambda_{u\to w}} P(o[u]\lambda_{\mathrm{out}(u)}|\lambda_{\mathrm{in}(u)}) \sum_{f(\lambda_{\mathrm{in}(w)})}\delta_{f(\lambda_{\mathrm{in}(w)}),o[w]}\, P(f(\lambda_{\mathrm{in}(w)})|\lambda_{\mathrm{in}(w)}) \\
&\quad= \sum_{\lambda_{u\to w}} P(o[u]\lambda_{\mathrm{out}(u)}|\lambda_{\mathrm{in}(u)})\cdot P(o[w]|\lambda_{\mathrm{in}(w)}) \\
&\quad= P(o[uw]\lambda_{\mathrm{out}(u)\setminus\{u\to w\}}|\lambda_{\mathrm{in}(u)\cup\mathrm{in}(w)\setminus\{u\to w\}}),
\end{aligned}$$

where the step $(*)$ uses the fact that upon regarding f as a list of elements of O_w indexed by $X_{\mathrm{in}(w)}$, all components of this list except for $f(\lambda_{\mathrm{in}(w)})$ simply drop out by normalization of probability.

This finishes our construction of making the information processing at w deterministic by pushing back the randomness needed u along the edge $u \to w$. The induction step can now be

completed as in the proof of Theorem 4.4: the hidden variables $\lambda_{\mathrm{in}(w)}$ have to be regarded as part of the outcomes $o[v]$ for those nodes v with $v \to w$, so that we obtain an induced classical correlation on the restricted graph $G\backslash\{w\}$. Applying the induction hypothesis to this situation then results in a classical model for the corresponding correlation $P(\tilde{o}[V\backslash\{w\}]) = P(o[V\backslash\{w\}]\lambda_{\mathrm{in}(w)})$ with the property that all information processing at the non-root nodes is deterministic. Putting this together with the deterministic processing P' at node w completes the induction step. □

By virtue of Theorem 4.4, this also shows that any classical correlation can be arbitrarily well approximated by one in which all conditional distributions $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$ are deterministic. We suspect that Conjecture 4.5 also holds without the approximation, but we do not have a proof of this.[2] The upcoming Theorem 7.3 might be helpful for achieving a complete proof.

[2]Our partial proof could possibly be extended to a full proof if the category of measurable spaces was cartesian closed, but this seems to be an open problem. See http://mathoverflow.net/questions/1388/is-there-a-natural-measures-on-the-space-of-measurable-functions.

5. Correlations within any theoretical framework

It should be useful to have a definition of correlation not just for the classical case, but also for other mathematical frameworks for physical theories. In particular, there should be definitions of quantum correlation and of general probabilistic correlations. We consider this problem by outlining the form in which a framework for physical theories, like "classical", "quantum" or "general probabilistic", needs to be given in order for it to enable a corresponding definition of correlation. Our answer to this turns out to be very similar to the definition of **causal operational-probabilistic theory** of Chiribella, D'Ariano and Perinotti [29]. The reason for tackling this very general problem at this stage is that it seems to give the most convenient rigorous definition of quantum correlation. Those readers who are not interested in mathematically precise definitions can directly proceed to the informal definition of quantum correlations in Section 6. Considerations similar to the ones of this section can also be found in the independent work of Henson, Lal and Pusey [16].

So what pieces of structure did we need in the previous section to define classical correlations? For one thing, each edge of the graph $G = (V, E)$ carries a classical system defined in terms of a measurable space. Second, the nodes $v \in V$ in turn are implemented as "information processing gates" given by stochastic instruments indexed by outcomes $o[v]$; here, an individual operation is typically subnormalized, i.e. happens with a state-dependent success probability. Only when taking the sum over all $o[v]$ do we get a normalization-preserving operation.

In general, the dichotomy between labelling the edges by state spaces and the nodes by operations suggests that we should assume that a framework for physical theories should be given by a **category** which we denote $\mathtt{C}$.[3] Then, each edge gets labelled by an object in $\mathtt{C}$, while each node gets labelled by a morphism of $\mathtt{C}$. More precisely, we assume that $\mathtt{C}$ is a (strict) **symmetric monoidal category** whose morphisms play the role of the operations; in particular, this means that no assumption like normalization or preservation of probability should be made when defining the morphisms of $\mathtt{C}$. In order to be able to make sense of the requirement that the every node carries one operation *for each outcome* and, when summing over all outcomes, the resulting total operation is normalized (more accurately, normalization-preserving), we need to make several additional assumptions:

(cat.i) Every morphism set $\mathtt{C}(X, Y)$ is a convex cone: any morphism in $\mathtt{C}(X, Y)$ can be multiplied by a positive scalar, any two morphisms $\mathtt{C}(X, Y)$ can be added, and there is a neutral element $0 \in \mathtt{C}(X, Y)$. More precisely, $\mathtt{C}(X, Y)$ is assumed to be a module over the semiring $\mathbb{R}_{\geq 0}$. Intuitively, the consequences of this are that we can take stochastic mixtures of operations, and the normalization of a morphism can be rescaled by simply multiplying the morphism by a scalar.

(cat.ii) The composition of morphisms

$$\mathtt{C}(X, Y) \times \mathtt{C}(Y, Z) \to \mathtt{C}(X, Z)$$

is bilinear with respect to this $\mathbb{R}_{\geq 0}$-module structure. In other words, composition distributes both over addition and scalar multiplication of morphisms. Likewise, the tensor product

$$\mathtt{C}(X_1, Y_1) \times \mathtt{C}(X_2, Y_2) \to \mathtt{C}(X_1 \otimes X_2, Y_1 \otimes Y_2)$$

is also bilinear.

[3] We refer to [30, 31] for background on the relevant pieces of category theory.

(cat.iii) In addition, there is a distinguished symmetric monoidal subcategory $\mathtt{C}_1 \subseteq \mathtt{C}$ containing all identity morphisms. We think of $\mathtt{C}_1$ as the subcategory which contains all the normalized operations.

(cat.iv) We have $\mathtt{C}(I,I) = \mathbb{R}_{\geq 0}$ and $\mathtt{C}_1(I,I) = \{1\}$, where I is the monoidal unit of $\mathtt{C}$ (and $\mathtt{C}_1$). In other words, there is exactly one normalized operation which does nothing to nothing, and this operation is denoted 1. The operations which act on nothing are precisely the scalar multiples of 1. This how we are going to obtain probabilities: as morphisms of type $I \to I$ in $\mathtt{C}$, which are nonnegative real numbers.

(cat.v) We have $\mathtt{C}_1(X,I) = \{1\}$, i.e. the unit object I is terminal in $\mathtt{C}_1$. In other words, every system X can be discarded in a unique way in the sense that there is exactly one normalized operation from X to no system which we denote $\tau_X : X \to I$ and draw as "$\underline{=}$". This assumption corresponds to the causality assumption of [29], which also forms a central aspect of the "causal categories" of [32], where the use of the ground symbol "$\underline{=}$" for τ has been introduced.

This finishes the list of requirements that the category $\mathtt{C}$ should satisfy in order for it to be possible to define the notion of "$\mathtt{C}$-correlation". For example, in the quantum case that we will consider in the next section, $\mathtt{C}$ will be the category of completely positive maps in which Hilbert spaces are the objects and completely positive maps between their sets of states will be the morphisms. As a more basic example, we can also define the classical correlations from the previous section like this by defining a category in which the objects are measurable spaces and the morphisms are stochastic operations. We will do this in more detail below.

Once a category $\mathtt{C}$ together with $\mathtt{C}_1 \subseteq \mathtt{C}$ as above has been determined, how do we obtain the definition of $\mathtt{C}$-correlation on a graph G? The basic idea is to interpret G as a diagram within the category $\mathtt{C}$, assuming that G has been labelled by an object $X_e \in \mathtt{C}$ for every edge $e \in E$, and by a morphism

$$g_v \in \mathtt{C}\left(\bigotimes_{e \in \mathrm{in}(v)} X_e, \bigotimes_{e \in \mathrm{out}(v)} X_e\right)$$

for every node $v \in V$. Labelling the edges by objects and the nodes by morphisms is essentially the same idea as the one which underlies "quantum picturalism" [33], namely the graphical calculus for symmetric monoidal categories [34]. This is another reason for why we believe that taking the hidden variables to live on the edges—as opposed to the nodes—is the appropriate thing to do. Once G can be interpreted as a diagram like this, it evaluates to an overall composite morphism $I \to I$, which (cat.iv) guarantees to be a nonnegative number. We interpret this number as the probability that the specific labelled graph G "happens", exactly as in Hardy's operator-tensor formulation [12].

More precisely, we want to label each node v not just by a single operation, but by a whole collection of operations indexed by the outcome $o[v]$. In other words, we need to assume that v is labelled by a function from outcomes to morphisms,

$$f_v : O_v \longrightarrow \mathtt{C}\left(\bigotimes_{e \in \mathrm{in}(v)} X_e, \bigotimes_{e \in \mathrm{out}(v)} X_e\right). \tag{5.1}$$

This function is assumed to assign to every outcome $o[v]$ a $\mathtt{C}$-operation $f_v(o[v]) : \bigotimes_{e\in \mathrm{in}(v)} X_e \to \bigotimes_{e\in \mathrm{out}(v)} X_e$ in such a way that $\sum_{o[v]} f_v(o[v])$ is normalized, i.e. lies in the subcategory $\mathtt{C}_1$. We can call this f_v a **$\mathtt{C}$-instrument** in analogy with the notion of quantum instrument, since it is a collection of operations indexed by an outcome such that the total operation obtained by marginalization (summing) over the outcome is normalized. The $\mathtt{C}$-instrument f_v plays the role of the "information processing gate" turning all the incoming systems into outgoing systems and producing a classical outcome while doing so. We think of a particular operation $f_v(o[v])$ as realized at node v exactly when the particular outcome $o[v]$ occurs. So when computing the probability $P(o[V])$ of a specific joint outcome $o[V]$ to occur, every node $v \in V$ is therefore labelled by exactly one morphism $f_v(o[v])$. Then one can try to interpret the graph G, labelled by the objects X_e and these particular morphisms $f_v(o[v])$, as a diagram in $\mathtt{C}$, which evaluates to an element of $\mathtt{C}(I, I)$. Indeed, the graphical calculus for symmetric monoidal categories, as formalized by Joyal and Street [34, Sec. 2.1] as the definition of "value" of a progressive polarized diagram, should be the appropriate way of doing this. By (cat.iv), this value $I \to I$ is a nonnegative real number, which we take to be the definition of $P(o[V])$.

However, there is an annoying technical problem with the idea of interpreting the graph G, when appropriately labelled, as a diagram in $\mathtt{C}$. Strictly speaking, the two tensor products $\bigotimes_{e\in \mathrm{in}(v)} X_e$ and $\bigotimes_{e\in \mathrm{out}(v)} X_e$ are not defined in $\mathtt{C}$: as a symmetric monoidal category, we can only make sense of binary tensor products in $\mathtt{C}$.[4] Hence, following Joyal and Street [34], we equip each set $\mathrm{in}(v)$ and $\mathrm{out}(v)$ with a particular *linear ordering*. We will show later that the resulting set of $\mathtt{C}$-correlations does not depend on this choice of ordering. Having fixed such orderings, we can now make sense of the tensor products $\bigotimes_{e\in \mathrm{in}(v)} X_e$ and $\bigotimes_{e\in \mathrm{out}(v)} X_e$ by defining them as iterated binary tensor products with respect to the given ordering.

In order to describe the notion of value in more detail, we need to introduce some additional terminology. The details of this will differ from the treatment of Joyal and Street, but it should be clear to the interested reader how to translate their definitions into ours. The general notion of value applies to directed acyclic graphs which may have incoming and outgoing "half-edges", i.e. edges which have only a source or target node within the graph. We can formalize this by defining a **directed acyclic graph with boundary** G to consist of a set V of nodes, a set E of ordinary edges, a set E_{in} of incoming edges, and a set E_{out} of outgoing edges, together with source and target maps

$$s : E \cup E_{\mathrm{out}} \to V, \qquad t : E \cup E_{\mathrm{out}} \to V.$$

So if we put $\mathrm{in}(v) := t^{-1}(v)$, then this set of incoming edges at v also includes all half-edges that have v as their target, and similarly for $\mathrm{out}(v)$. The acyclicity condition is then as before in Definition 2.2, and G is acyclic if and only if the restricted graph (V, E) without the half-edges is so. During the discussion of value, we use the term "graph" as a synonym of "directed acyclic graph with boundary". As a special case, the boundary may be empty, in which case the definition reduces to the usual notion of directed acyclic graph which we have been using until now.

[4] There is an "unbiased" definition of monoidal category in which tensor products are not necessarily binary [35, Sec. 3.1], but an ordering of the factors does still need to be chosen. More generally, there is also a definition of "fat symmetric multicategory" [35, App. A], which is more general than the definition of "order-independent unbiased symmetric monoidal category" that we would need. However, there does not yet seem to be any theorem on interpreting directed acyclic graphs as diagrams in symmetric monoidal categories, which is the reason for the technical complications that we are facing here. See also nforum.mathforge.org/discussion/3101/symmetric-monoidal-category/?Focus=46303#Comment_46303.

Such a graph is *polarized* if each set $\mathrm{in}(v)$ and $\mathrm{out}(v)$ is equipped with a linear ordering. It is *anchored* if both E_{in} and E_{out} are equipped with linear orderings. A *valuation* in $\mathtt{C}$ is a labelling of E, E_{in} and E_{out} by objects $X_e \in \mathtt{C}$ and of V by morphisms g_v of $\mathtt{C}$ such that each f_v is of type

$$g_v \;:\; \bigotimes_{e \in \mathrm{in}(v)} X_e \longrightarrow \bigotimes_{e \in \mathrm{out}(v)} X_e,$$

where it is understood that the tensor products are taken with respect to the chosen orderings.

Then Joyal and Street [**34**, Sec. 2.1] show that it is possible to assign a value to this graph in the sense of regarding it as a diagram in $\mathtt{C}$ and composing it to a morphism

$$\nu(G) : \bigotimes_{e \in E_{\mathrm{in}}} X_e \longrightarrow \bigotimes_{e \in E_{\mathrm{out}}} X_e,$$

where again these tensor products are with respect to the chosen orderings on E_{in} and E_{out}. If $E_{\mathrm{in}} = E_{\mathrm{out}} = \emptyset$, as in the case of interest to us, then $\nu(G) : I \to I$. In fact, it can be shown that there is a unique way of assigning a value satisfying the following properties:

(val.i) If G is a graph containing exactly one node, then $\nu(G)$ coincides with the morphism labelling that node.

(val.ii) If G_1 and G_2 are graphs equipped with valuations in $\mathtt{C}$, then also the disjoint union graph $G_1 + G_2$ carries a valuation in $\mathtt{C}$, and

$$\nu(G_1 + G_2) = \nu(G_1) \otimes \nu(G_2).$$

Intuitively, this corresponds to placing the two diagrams G_1 and G_2 side by side, and hence their value should be interpreted as the tensor product of the individual values.

(val.iii) The second requirement is similar and refers to composition of graphs: if G_1 and G_2 are such that there is an equality of ordered sets $E_{1,\mathrm{out}} = E_{2,\mathrm{in}}$, and both G_1 and G_2 carry a valuation such that the object assigned to a half-edge $e \in E_{1,\mathrm{out}}$ coincides with the object associated to the corresponding half-edge $e \in E_{2,\mathrm{in}}$, then the two graphs can be composed to $G_2 \circ G_1$ in the obvious way, and this composite graph also carries a valuation. The requirement now is that

$$\nu(G_2 \circ G_1) = \nu(G_2) \circ \nu(G_1),$$

where "$\circ$" on the right-hand side is composition in $\mathtt{C}$.

(val.iv) Taking the value is a linear function of the morphism g_v at each node $v \in V$.

Armed with these preliminary considerations, we can now get to our main result of this section. If we have a graph without boundary G, and G is labelled by objects on the edges and $\mathtt{C}$-instruments (5.1) on the nodes, then every joint outcome $o[V]$ can be assigned a number $P(o[V])$ by evaluating it as a diagram in $\mathtt{C}$ as described above. The main point is this:

5.1. Theorem. *The resulting numbers $P(o[V])$ form a correlation on G.*

See also [**36**] for more general considerations relating assumption (cat.v) to the no-signaling property.

Proof. This proof generalizes the one of Proposition 4.2. In particular, the overall proof strategy is the same, and we will see that Proposition 4.2 can indeed be regarded as an instance of the present claim.

We start by deriving the following generalization of (4.3): if $W \subseteq V$ is a set of nodes with $\mathrm{pst}(W) = W$, then the marginal $P(o[W])$ can be computed using the following steps:

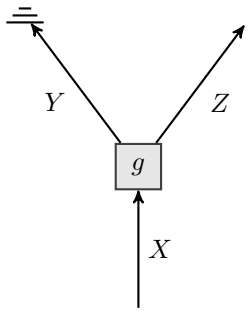

FIGURE 7. Taking the marginal of a morphism in $\mathtt{C}$ over an output system Y is defined in terms of composition with the unique morphism $\tau_Y \in \mathtt{C}_1(Y, I)$ drawn as "$\underline{=}$". In the quantum case, this is an application of partial trace over Y.

- Take the restricted graph $(W, E|_W)$, where $E|_W$ contains exactly those edges whose source and target both lie in W. Just as the original G, this is also a graph without boundary.
- Equip $(W, E|_W)$ with the same objects X_e as edge labels, and label the nodes by the $\mathtt{C}$-instruments whose components are the corresponding marginal operations (Figure 7). Here, the marginal of an operation $f_v(o[v])$ needs to be taken over all those edges $e \in \mathrm{out}(v)$ whose target $\mathrm{tar}(e)$ does not lie in W. It follows from bilinearity (cat.ii) and the fact that the marginalization morphisms $\tau_{\mathrm{tar}(e)}$ are in the subcategory $\mathtt{C}_1$ that the new marginal operations do also form a $\mathtt{C}$-instrument.
- Take the linear orderings on the new sets $\mathrm{in}(v)$ and $\mathrm{out}(v)$ for all $v \in W$ to be those induced from G.
- Evaluate this graph to a family of numbers $P(o[W])$ by applying ν as described above.

We prove that the resulting numbers $P(o[W])$ coincide with the corresponding marginal of $P(o[V])$ by induction on the size of $V \backslash W$. The base case is $W = V$, for which the claim is simply the definition of $P(o[V])$. For the induction step, we need to prove that the claim holds for a given W with $\mathrm{pst}(W) = W$, while assuming that it holds for all bigger W' in place of W which likewise satisfy $\mathrm{pst}(W') = W'$. For the given W, we can choose a node w with $\mathrm{pa}(w) \subseteq W$, as in the proof of (4.3). Upon putting $W' := W \cup \{w\}$, it is then guaranteed that $\mathrm{pst}(W') = W'$. We therefore can apply the induction assumption, which gives that the marginal $P(o[W'])$ can be computed by applying the above steps to W'; this is depicted in Figure 8(a). Then it remains to evaluate the sum over $o[w]$, in order to show that the further restricted marginal

$$P(o[W]) = \sum_{o[w]} P(o[W']) \tag{5.2}$$

can likewise be computed by following the above prescription. In order to prove this, we use the linearity property (val.iv) to interpret the right-hand side as the value of the diagram arising from the valuation of $(W', E|_{W'})$ by objects and operations such that each $v \in W$ is labelled by the morphism $f_v(o[v])$ where the value $o[v]$ is the given one on the left-hand side of (5.2), while the "new" node w is labelled by $\sum_{o[w]} f_w(o[w])$. Because the latter is a morphism in $\mathtt{C}_1$, we can apply assumption (cat.v): since I is a terminal object in $\mathtt{C}_1$, the set $\mathtt{C}_1\big(\bigotimes_{e \in \mathrm{in}(w)} X_e, I\big)$ has exactly one element. In particular, the elements $\sum_{o[w]} f_w(o[w]) \in \mathtt{C}_1\big(\bigotimes_{e \in \mathrm{in}(w)} X_e, I\big)$ and

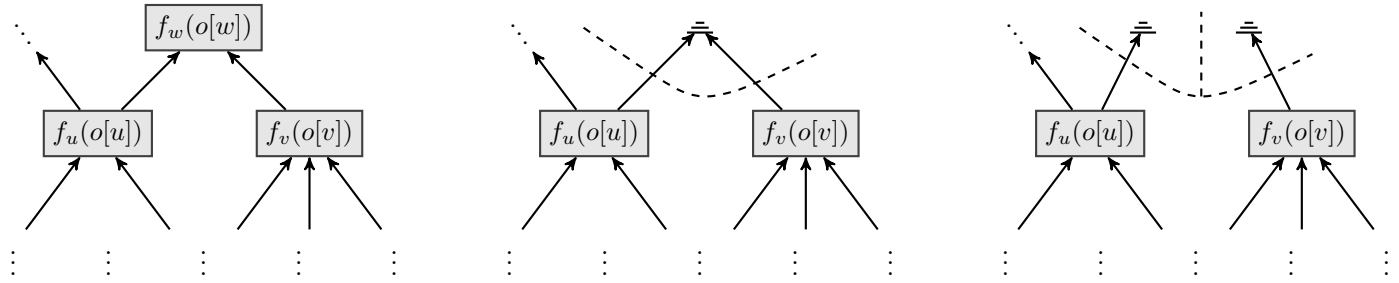

(a) (b) (c)

FIGURE 8. Illustration of the induction step in the proof of Theorem 5.1 for part of an example graph. The ground symbol $\overline{=}$ represents the unique normalized operation $\tau_X \in \mathtt{C}_1(X, I)$ from any system X to no system.

$\bigotimes_{e\in\mathrm{in}(w)} \tau_{X_e} \in \mathtt{C}_1\big(\bigotimes_{e\in\mathrm{in}(w)} X_e, I\big)$ are equal,

$$\sum_{o[w]} f_w(o[w]) = \bigotimes_{e\in\mathrm{in}(w)} \tau_{X_e}, \tag{5.3}$$

where τ_{X_e} is itself the unique element of $\mathtt{C}_1(X_e, I)$. Intuitively, the idea is that $\sum_{o[w]} f_w(o[w])$ is the unique normalized operation going from the composite system $\bigotimes_{e\in\mathrm{in}(w)} X_e$ to no system, as illustrated in Figure 8(b), and hence coincides with throwing away all systems in parallel, which is given by the morphism $\bigotimes_{e\in\mathrm{in}(w)} \tau_{X_e}$ as depicted in Figure 8(c). Now we can regard the graph $(W', E|_{W'})$ as the composition of two parts as indicated by the dashed line in Figure 8(b). The upper part consists of only one node with ingoing edges and is labelled by the morphism (5.3), and hence by (val.i) its value coincides with this very morphism. Since this morphism itself factors into a tensor product, rule (val.ii) tells us that we obtain the same value if we furnish each incoming edge with its own target node, as in Figure 8(c). Then again by the rules (val.ii) and (val.iii), the overall diagram value coincides with the one obtained by taking $(W, E|_W)$ and labelling it with marginal morphisms as described above. This completes the induction step.

Now we can make use of the above prescription. First, if $W = \emptyset$, then we obtain the marginal $P(o[\emptyset])$, which simply is a number corresponding to the normalization of the distribution $P(o[V])$, which should be 1. What the above prescription tells is that we can compute $P(o[\emptyset])$ as the value of the empty graph, which is $\mathrm{id}_I \in \mathtt{C}(I, I)$. As a number, this is indeed 1.

Second, we can also show that the "binary" correlation equation (3.3) with the assumption $\mathrm{pst}(U) = U$ and $\mathrm{pst}(W) = W$ holds by applying the above prescription to the subgraph $(U \cup W, E|_{U\cup W})$. Since this subgraph decomposes into the disjoint union of $(U, E|_U)$ and $(W, E|_W)$, the claim now follows immediately from (val.ii). □

So for any category of operations $\mathtt{C}$, we now have a precise definition of the set of $\mathtt{C}$-correlations. However, this definition was phrased in terms of the chosen orderings on each $\mathrm{in}(v)$ and $\mathrm{out}(v)$; since these are not part of the data of the original graph G, we still need to show that the set of $\mathtt{C}$-correlations is independent of that choice.

5.2. PROPOSITION. *The set of $\mathtt{C}$-correlations on G does not depend on the particular orderings chosen on the sets* $\mathrm{in}(v)$ *and* $\mathrm{out}(v)$.

SKETCH OF PROOF. It is enough to show that the ordering on $\mathrm{in}(v)$ and $\mathrm{out}(v)$ for some fixed v can be modified arbitrarily without changing the set of $\mathtt{C}$-correlations. Since any two orderings are related by a permutation, and any permutation can be generated from adjacent transpositions which do nothing but exchange a pair of neighboring edges, it is enough to consider the case that only two neighboring edges e_1, e_2 are exchanged at some fixed node v, either among the incoming edges or among the outgoing edges. We then need to show that if $P(o[V])$ is a $\mathtt{C}$-correlation with respect to the original ordering, then it is also a $\mathtt{C}$-correlation with respect to the new ordering. Since $P(o[V])$ is a $\mathtt{C}$-correlation with respect to the original ordering by assumption, we know that there exists a labelling of edges by objects X_e of $\mathtt{C}$ and of nodes by instruments f_v in $\mathtt{C}$ which reproduces the probabilities. Now we can construct a labelling of the graph with the new ordering by replacing the instrument f_v with the instrument given by the composition of each component operation $f_v(o[v])$ with the symmetry (swap) $X_{e_1} \otimes X_{e_2} \longrightarrow X_{e_2} \otimes X_{e_1}$, while retaining all other labels. If e_1 and e_2 are incoming edges of v, then the symmetry has to be applied before f_v; otherwise, if $e_1, e_2 \in \mathrm{out}(v)$, it should be used after. This defines a new set of valuations of the same graph with the new ordering which reproduces the same probabilities $P(o[V])$ thanks to the invariance properties of the graphical calculus as given in terms of the map ν [**34**, Sec. 2.1]. □

We would like to reemphasize that having to deal with these orderings explicitly is not a shortcoming of the use of category theory per se; it is rather due to the fact that the appropriate pieces of category theory do not seem to have been developed *yet*. Recent work like [**37**] paves the way towards improving this situation.

If one follows the "hidden variables everywhere" paradigm advocated in this paper, then there should be no difference between the causal structure corresponding to direct causal links $v \to w$ and the one corresponding to indirect causal links $v \rightsquigarrow w$. The intuitive reason is this: clearly, every correlation explainable in terms of hidden variables on the given causal structure G is likewise explainable by the enlarged causal structure comprising all direct *and indirect* causal links. To see that the converse also holds, one can simulate the propagation of a hidden variable along an indirect causal link by passing it *incognito* along all the direct causal links making up the indirect link as additional information. So there is ultimately no reason to distinguish between a direct link $v \to w$ and an indirect link $v \rightsquigarrow w$, and so it should be sufficient to define the causal structure only in terms of $v \rightsquigarrow w$.

5.3. THEOREM. *The set of* $\mathtt{C}$*-correlations depends only on the partial order (=causal set) induced by* G*, as defined in Section 2.*

SKETCH OF PROOF. It is sufficient to show the following: if $u, v, w \in V$ are nodes such that $u \to v$ and $v \to w$ but $u \not\to w$, then the graph $G' := (V, E \cup \{u \to w\})$ differs from G only in that the edge $u \to w$ has been added. It is enough to show that G' has the same set of $\mathtt{C}$-correlations as G, since applying this statement repeatedly proves the claim.

As noted above, every $\mathtt{C}$-correlation on G is clearly also a $\mathtt{C}$-correlation on G': any labelling of G by objects and instruments in $\mathtt{C}$ can be extended to a labelling of G' by assigning the monoidal unit $I \in \mathtt{C}$ to the new edge $u \to w$, while retaining all other labels. Since labelling an edge by I is the same as not considering that edge at all, the resulting correlation is the same.

Conversely, suppose that $P(o[V])$ is a $\mathtt{C}$-correlation on G' constructed in terms of some labelling by objects $(X'_e)_{e \in E \cup \{u \to w\}}$ and instruments $(f'_x)_{x \in V}$. We can turn this into a corresponding labelling of G by again retaining all labels, except for the replacements

$$X_{u \to v} := X'_{u \to v} \otimes X'_{u \to w}, \qquad X_{v \to w} := X'_{v \to w} \otimes X'_{u \to w}, \qquad f_v(o[v]) := f'_v(o[v]) \otimes \mathrm{id}_{X'_{u \to w}},$$

where we have assumed that the additional edge $u \to w$ comes last in the ordering of $\mathrm{in}(v)$ and $\mathrm{out}(v)$ on G'. Intuitively, the idea is this: in G', the information on $X_{u\to w}$ gets passed directly from u to w, while G needs to use the intermediate node v as a relay for sending the information from u to w via v. By considering the value map ν only on the subgraph with boundary induced by the three nodes u, v and w, the properties (val.i)–(val.iii) imply that the resulting number $P(o[V])$ ends up being the same. □

We now illustrate this categorical formalism with the example of classical correlations, before applying it to quantum correlations in the next section. Classical correlations are defined in terms of a category $\mathtt{StOp}$ whose objects are measurable spaces (X, Σ_X), and whose morphisms $(X, \Sigma_X) \to (Y, \Sigma_Y)$ are stochastic operations; see Appendix A. The monoidal product $\otimes$ of measurable spaces is simply given by their cartesian product. The unit object I is any fixed measurable space having exactly one point. The subcategory of normalized operations $\mathtt{StOp}_1 \subseteq \mathtt{StOp}$ contains precisely those stochastic operations $P(y|x)$ which are normalized, i.e. $\int_y P(y|x) = 1$ for all $x \in X$. It is straightforward to check that this turns $\mathtt{StOp}$ into a symmetric monoidal category satisfying our assumptions (cat.i)–(cat.v).

Instantiating the above general definition for the category $\mathtt{StOp}$ results in the notion of $\mathtt{StOp}$-correlation, which coincides exactly with the notion of classical correlation from Definition 4.1. The integrals in (4.2) arise from the composition of stochastic operations as defined in (A.1).

Besides the definition of quantum correlations in the next section, we also suspect that it should be possible to define general probabilistic [38] correlations as arising from a suitable category in the analogous way. Currently, the problem with defining general probabilistic theories is that a corresponding category of general probabilistic operations does not yet seem to have been worked out. However, it has been shown in independent work of Henson, Lal and Pusey [16] that there exist correlations in certain scenarios that are not general probabilistic correlations. This should be seen in contrast to the case of Bell scenarios, where the correlations correspond to the no-signaling boxes (Theorem 3.3), and it is well-known that every no-signaling box arises from the general probabilistic theory known as "box world" [39].

Sometimes a certain framework for physical theories is clearly more powerful than another one. For example, it is commonly believed that everything that can be done classically can also be done quantumly, and this should imply that every classical correlation is also a quantum correlation. One can try to prove statements of this form by making use of this result:

5.4. Proposition. *Let* $\mathtt{C}$ *and* $\mathtt{D}$ *be symmetric monoidal category of operations satisfying (cat.i)–(cat.v). If there is a linear symmetric monoidal functor* $F : \mathtt{C} \to \mathtt{D}$ *with* $F(\mathtt{C}_1) \subseteq \mathtt{D}_1$*, then every* $\mathtt{C}$*-correlation is also a* $\mathtt{D}$*-correlation.*

Here, linearity of F means that F preserves addition and scalar multiplication of morphisms. The functor is monoidal if it comes equipped with isomorphisms $F(X \otimes_{\mathtt{C}} Y) \cong F(X) \otimes_{\mathtt{D}} F(Y)$ satisfying the appropriate coherence laws, one of which involves the symmetry isomorphisms. These coherence requirements are technical conditions guaranteed to hold in any naturally arising situation.

Sketch of Proof. Let $P(o[V])$ be a $\mathtt{C}$-correlation arising via objects $(X_e)_{e\in E}$ and instruments $(f_v)_{v\in V}$. Linearity of the functor F and $F(\mathtt{C}_1) \subseteq \mathtt{D}_1$ imply that the instrument components $f_v(o[v])$ map to $\mathtt{D}$-operations $F(f_v(o[v]))$ which form an instrument themselves that we denote $F(f_v)$. Hence we obtain a labelling of G in $\mathtt{D}$ with objects $(F(X_e))_{e\in E}$ and instruments $(F(f_v))_{v\in V}$. The properties (val.i)–(val.iii) then guarantee that the resulting probabilities $P(o[V])$ are the same, since they

imply that applying ν commutes with F, and F maps the numbers in $\mathtt{C}(I, I)$ to the same numbers in $\mathtt{D}(I, I)$ thanks to linearity. $\square$

6. Quantum correlations

Our world is governed, at least on microscopic scales, by the laws of quantum mechanics and quantum field theory. It was the revolutionary insight of Bell [1] that this leads to correlations (on the graph of Figure 1) which are not classical. As a consequence, either our world cannot be described in terms of classical probability theory, or the given causal structure is not actually the correct one describing the given correlations. The first option is generally known as **non-realism** [40, 41], while the second option is referred to as **non-locality**, in the sense of the existence of non-local interactions[5]. On the mathematical level, the non-realism solution is preferable since non-locality requires fine-tuning to explain why these non-local interactions cannot be exploited for instantaneous signalling [15].

Using the definitions and results from the previous section, we will now define what a **quantum correlation** is in our more general formalism. The basic idea is that instead of carrying measurable spaces as in the classical case, the edges should carry the quantum analogue of those, i.e. Hilbert spaces. Similarly, instead of having stochastic operations at the nodes, the nodes should now be labelled by **quantum operations**. Roughly speaking, quantum correlations are exactly those joint distributions $P(o[V])$ that one can obtain upon interpreting the given graph in Hardy's operator-tensor formulation of quantum theory [12]. However, since the operator-tensor formulation has so far only been worked out for finite-dimensional quantum systems, we cannot use it for our definition of quantum correlation. We prefer to be general and allow systems of arbitrary finite or infinite dimension, since we have no way to tell how much stronger quantum correlations based on infinite-dimensional Hilbert space may be. While we suspect that there should be a quantum analogue of Theorem 4.4, already stating and proving such a theorem requires having a general definition. Without a general definition, there is no way to know whether quantum correlations with infinite-dimensional Hilbert spaces can be stronger than those based on finite dimensions or not.

So in order to apply the definitions and results from the previous section, we need to explain how quantum operations form a symmetric monoidal category, which we denote $\mathtt{QOp}$. This works as follows: the objects of $\mathtt{QOp}$ are Hilbert spaces $\mathcal{H}$ of arbitrary dimension, finite or infinite. These can be tensored in the usual way; we refer to [42, 43] for background on the definition of this tensor product in the infinite-dimensional case. To every Hilbert space $\mathcal{H}$ we associate its space of trace class operators $\mathcal{S}(\mathcal{H})$. One can understand the letter "$\mathcal{S}$" as suggestive of either "state" or of "Schatten class" [44]. The quantum states (density matrices) on $\mathcal{H}$ are precisely those positive elements of $\mathcal{S}(\mathcal{H})$ that have unit trace. With this in mind, we can now define the morphisms of $\mathtt{QOp}$: for Hilbert spaces $\mathcal{H}$ and $\mathcal{K}$, the morphisms forming the set $\mathtt{QOp}(\mathcal{H},\mathcal{K})$ are defined to be the **completely positive** maps $\mathcal{S}(\mathcal{H}) \to \mathcal{S}(\mathcal{K})$. Note that our usage of the term "quantum operation" differs from the usual one [45, Sec. 8.2.4] in that we do not impose subnormalization. The normalized maps $\mathcal{S}(\mathcal{H}) \to \mathcal{S}(\mathcal{K})$ forming the subcategory $\mathtt{QOp}_1$ are those quantum operations that are trace-preserving, i.e. the quantum channels. It is then straightforward to check that our requirements (cat.i)–(cat.v) are satisfied.

6.1. Definition. *$P(o[V])$ is a **quantum correlation** if it is a $\mathtt{QOp}$-correlation.*

We refer to the proof of Proposition 6.5 for an example of how to unfold this definition for a concrete graph. Thanks to the general Theorem 5.1 from the previous section, we obtain immediately:

[5] Unfortunately, the term *nonlocality* is also often used in a much broader sense, as a synonym for the existence of non-classical correlations [2].

6.2. Proposition. *If $P(o[V])$ is a quantum correlation, then it is a correlation.*

Moreover, we also would like to apply Proposition 5.4 in order to show that every classical correlation is also a quantum correlation. However, the problem with this is that we do not know how to construct any functor $\mathtt{StOp} \to \mathtt{QOp}$ satisfying our requirements, since there is no obvious way to assign a Hilbert space to any given measurable space (X, Σ). The reader not interested in the reasons for this may skip the following remark.

6.3. Remark. There is a standard construction assigning a Hilbert space $L^2(X, \Sigma, P)$ to any *measure space* (X, Σ, P), but this Hilbert space depends crucially on the measure P, so that it is not applicable to a mere measurable space (X, Σ). In fact, even for the full subcategory $\mathtt{FinStOp} \subseteq \mathtt{StOp}$ containing only the finite measurable spaces $(X, \mathbf{2}^X)$, we have not been able to find any functor $F : \mathtt{FinStOp} \to \mathtt{QOp}$ that would satisfy the assumptions of Proposition 5.4. Any obvious candidate should be given on objects by assigning to every finite set X the Hilbert space $\mathbb{C}^X$, which has a canonical orthonormal basis given by $\{|x\rangle\}_{x \in X}$. To any stochastic operation $P(y|x) : X \to Y$, we can try to assign the quantum operation

$$\mathcal{S}(\mathbb{C}^X) \to \mathcal{S}(\mathbb{C}^Y), \qquad |x_1\rangle\langle x_2| \mapsto \sum_{y_1, y_2} \sqrt{P(y_1|x_1)\, P(y_2|x_2)}\, |y_1\rangle\langle y_2|, \tag{6.1}$$

which can be understood as being represented by the single Kraus operator

$$\mathbb{C}^X \to \mathbb{C}^Y, \qquad |x\rangle \mapsto \sum_y \sqrt{P(y|x)}\, |y\rangle.$$

It is straightforward to check that this assignment of a quantum operation to any stochastic operation is indeed a monoidal functor. However, due to the square root in (6.1), it fails to be linear—the sum of two classical operations will in general not map to the sum of the associated quantum operations! So, alternatively, one might try and come up with an alternative to (6.1) which is linear in P. The obvious candidate is a "decoherence in the canonical basis" approach, in which we would assign to $P(y|x)$ the quantum operation

$$\mathcal{S}(\mathbb{C}^X) \to \mathcal{S}(\mathbb{C}^Y), \qquad |x_1\rangle\langle x_2| \mapsto \delta_{x_1, x_2} \sum_y P(y|x_1)\, |y\rangle\langle y|, \tag{6.2}$$

which indeed has the advantage of being linear in P. However, this assignment lacks functoriality: the identity stochastic operation $X \to X$, represented by $P(x'|x) = \delta_{x', x}$, gets mapped to the quantum operation $\mathcal{S}(\mathbb{C}^X) \to \mathcal{S}(\mathbb{C}^X)$ which decoheres any quantum state with respect to the computational basis; in other words, this assignment of a quantum operation to any stochastic operation does not map identities to identities, as any functor should, and hence (6.2) does not define a functor.

So while we strongly suspect that every classical correlation is indeed a quantum correlation, we do not have any proof of this seemingly basic statement!

6.4. Conjecture. *Every classical correlation is also a quantum correlation.*

All that we can say for sure is that every classical correlation which can be realized with finite hidden variable spaces is also a quantum correlation. Providing a proof sketch of this is what we do now. By Theorem 4.4, this is arbitrarily close to proving the full conjecture, but does not go all the way.

PARTIAL PROOF. As outlined above, we do not know how to apply Proposition 5.4 to this situation, but turning any stochastic operation between finite sets into a quantum operation as in (6.2) does indeed do the trick, despite not being functorial. So for each edge $e \in E$, we replace the hidden variable space X_e, assumed to be finite, by the Hilbert space $\mathbb{C}^{X_e}$. We also replace every stochastic instrument $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$ by the corresponding "decoherence in the canonical basis" quantum instrument

$$O_v \times \mathcal{S}\left(\bigotimes_{e\in\mathrm{in}(v)} \mathbb{C}^{X_e}\right) \longrightarrow \mathcal{S}\left(\bigotimes_{e\in\mathrm{out}(v)} \mathbb{C}^{X_e}\right)$$

with same outcome set O_v. Then all states propagating along the edges will be diagonal in the canonical basis. This guarantees that the original joint probabilities $P(o[V])$ are recovered. □

In order to underline that our definition of quantum correlation is adequate, we compare it to the usual definition in the Bell scenario case:

6.5. PROPOSITION. *In the n-party Bell scenario, $P(o[a_1 \dots a_n x_1 \dots x_n s])$ is a quantum correlation if and only if it satisfies the free will equation (3.4) and the corresponding conditional probabilities admit a* ***quantum model****,*

$$P(o[a_1 \dots a_n]|o[x_1 \dots x_n s]) = \langle \psi_{o[s]} | E^{o[x_1]}_{o[a_1]} \otimes \dots \otimes E^{o[x_n]}_{o[a_n]} | \psi_{o[s]} \rangle \,, \tag{6.3}$$

where $\{|\psi_{o[s]}\rangle\}_{o[s]\in O_s}$ is a family of states on some composite Hilbert space $\mathcal{H}_1 \otimes \dots \otimes \mathcal{H}_n$ and the $E^{o[x_i]}_{o[a_i]}$ are positive operators on $\mathcal{H}_i$, not depending on $o[s]$, subject to the completeness relation $\sum_{o[a_i]} E^{o[x_i]}_{o[a_i]} = \mathbb{1}_{\mathcal{H}_i}$.

As in Proposition 4.3, we implicitly also assume that only those values of $o[x_i]$ and $o[s]$ are considered for which $P(o[x_i]) > 0$ and $P(o[s]) > 0$, since otherwise conditioning on these variables would not make sense.

PROOF. We begin by spelling out our definition of quantum correlation for the n-party Bell scenario in more detail. First, our definition requires that the edges of the graph are labelled by Hilbert spaces

$$\mathcal{H}_{x_i \to a_i}, \qquad \mathcal{H}_{s \to a_i}.$$

Second, there is a quantum instrument labelling the source node s. Since s does not have any incoming edges, this quantum instrument corresponds to a collection of subnormalized states indexed by the outcome $o[s] \in O_s$,

$$\rho_{o[s]} \in \mathcal{S}\left(\bigotimes_{i=1}^{n} \mathcal{H}_{s\to a_i}\right), \qquad \mathrm{tr}\left(\sum_{o[s]} \rho_{o[s]}\right) = 1,$$

where the prescription in the proof of Theorem 5.1 shows that $P(o[s]) = \mathrm{tr}(\rho_{o[s]})$. The same statement applies to the "choice of setting" nodes x_i, which likewise emit subnormalized states indexed by outcomes $o[x_i]$,

$$\sigma_{o[x_i]} \in \mathcal{S}(\mathcal{H}_{x_i\to a_i}), \qquad \mathrm{tr}\left(\sum_{o[x_i]} \sigma_{o[x_i]}\right) = 1,$$

and we similarly have $P(o[x_i]) = \mathrm{tr}(\sigma_{o[x_i]})$.

Finally, the quantum instrument labelling a "measurement" node a_i corresponds simply to a POVM[6],

$$M_{o[a_i]} \in \mathcal{B}\left(\mathcal{H}_{x_i \to a_i} \otimes \mathcal{H}_{s \to a_i}\right), \qquad \sum_{o[a_i]} M_{o[a_i]} = \mathbb{1}.$$

This data is related to the overall probabilities by

$$P(o[a_1 \ldots a_n x_1 \ldots x_n s]) = \mathrm{tr}\left[\left(\rho_{o[s]} \otimes \bigotimes_{i=1}^{n} \sigma_{o[x_i]}\right)\left(\bigotimes_{i=1}^{n} M_{o[a_i]}\right)\right], \tag{6.4}$$

where it is understood that the tensor factors in the tensor product Hilbert space carrying the "state" part need to be reordered so as to match those of the "measurement" part in order for the operator multiplication to make sense. (This is closely related to the ordering issues that we encountered in Section 5.)

We begin the proof with the "only if" part. One obtains the free will equation (3.4) by noting that Proposition 6.2 guarantees that $P(o[a_1 \ldots a_n x_1 \ldots x_n s])$ is a correlation, and therefore Theorem 3.3 applies. In order to obtain a representation of the form (6.3), we take the Hilbert spaces to be $\mathcal{H} := \mathcal{H}_{s \to a_i}$, on which we have positive operators

$$E^{o[x_i]}_{o[a_i]} := \frac{1}{P(o[x_i])} \mathrm{tr}_{\mathcal{H}_{x_i \to a_i}}\left(\left(\sigma_{o[x_i]} \otimes \mathbb{1}_{\mathcal{H}_{s \to a_i}}\right) M_{o[a_i]}\right). \tag{6.5}$$

Since $\sum_{o[a_i]} M_{o[a_i]} = \mathbb{1}$ and $\mathrm{tr}(\sigma_{o[x_i]}) = P(o[x_i])$, this indeed satisfies the completeness relation $\sum_{o[a_i]} E^{o[x_i]}_{o[a_i]} = \mathbb{1}_{\mathcal{H}_{x_i \to a_i}}$. The definition (6.5) formalizes the intuitive idea that $o[x_i]$ behaves like a choice of setting determining which POVM will be used at a_i. With this, we can rewrite (6.4) in the form

$$P(o[a_1 \ldots a_n x_1 \ldots x_n s]) = \mathrm{tr}\left[\left(\rho_{o[s]} \otimes \bigotimes_{i=1}^{n} \mathbb{1}_{\mathcal{H}_{s \to a_i}}\right)\left(\bigotimes_{i=1}^{n} E^{o[x_i]}_{o[a_i]}\right)\right] \cdot \prod_{i=1}^{n} P(o[x_i]),$$

where now the trace only ranges over the remaining Hilbert space $\bigotimes_{i=1}^{n} \mathcal{H}_{s \to a_i}$. If we now purify the states $\rho_{o[s]}$ jointly to unit vectors $|\psi_{o[s]}\rangle$ by adding an additional ancilla system, which should be regarded as becoming part of any one of the $\mathcal{H}_{s \to a_i}$, we can write this as

$$P(o[a_1 \ldots a_n x_1 \ldots x_n s]) = \left\langle \psi_{o[s]} \middle| E^{o[x_1]}_{o[a_1]} \otimes \ldots \otimes E^{o[x_n]}_{o[a_n]} \middle| \psi_s \right\rangle \cdot P(o[x_1]) \cdots P(o[x_n]) P(o[s]), \tag{6.6}$$

where we had to factor out an additional factor of $P(o[s])$, since $|\psi_{o[s]}\rangle$ is normalized, while $\mathrm{tr}(\rho_{o[s]}) = P(o[s])$. This gives the desired form (6.3) upon dividing by the free will equation (3.4).

Now for the "if" part. Suppose (3.4) and (6.3) hold; we can multiply these two equations and obtain (6.6). We rewrite this in the form (6.4) as follows. Let the Hilbert spaces on the edges be given by

$$\mathcal{H}_{s \to a_i} := \mathcal{H}_i, \qquad \mathcal{H}_{x_i \to a_i} := \mathbb{C}^{O_{x_i}},$$

where $\mathbb{C}^{O_{x_i}}$ is the Hilbert space with a canonical orthonormal basis $\{|o[x_i]\rangle\}$ indexed by the possible values of $o[x_i]$. Now define

$$\rho_{o[s]} := P(o[s]) \cdot |\psi_{o[s]}\rangle\langle\psi_{o[s]}|, \quad \sigma_{o[x_i]} := P(o[x_i]) \cdot |o[x_i]\rangle\langle o[x_i]|, \quad M_{o[a_i]} := \sum_{o[x_i]} |o[x_i]\rangle\langle o[x_i]| \otimes E^{o[x_i]}_{o[a_i]}.$$

[6]To begin with, the components of the quantum instrument at a_i are completely positive maps of type $\mathcal{S}(\mathcal{H}) \to \mathbb{C}$ for $\mathcal{H} = \mathcal{H}_{x_i \to a_i} \otimes \mathcal{H}_{s \to a_i}$, but it is well-known that these are in bijection with bounded positive operators on $\mathcal{H}$.

It is straightforward to check that these operators satisfy the properties mentioned above that are required for the operators to assemble into quantum instruments on the nodes s, x_i and a_i, and that (6.4) then indeed recovers the given correlation. □

The proofs of this proposition and the earlier Propositions 3.3 and 4.3 should make it clear how to translate back and forth between Bell scenarios in our formalism and the already existing vast literature on Bell scenarios, which uses the ordinary definitions (4.4) and (6.3).

We leave all further study of quantum correlations on arbitrary causal structure to future work. As a start, it should be verified that our general definition not only comprises the usual one for Bell scenarios, but also for all other concrete scenarios that have been studied until now [**8**, **9**, **10**, **11**], as far as precise definitions have been spelled out in these works. Then, it will be a challenging enterprise to find further examples of non-classical quantum correlations and try to answer the question in which scenarios non-classical quantum correlations exist and in which ones they do not; as a trivial example, the single-party Bell scenario does not have any non-classical quantum correlations. We believe that a lot remains to be explored, and works like the present one and [**10**, **16**] are but the beginning of a transformation of our understanding of quantum "nonlocality" and its applications to quantum information processing.

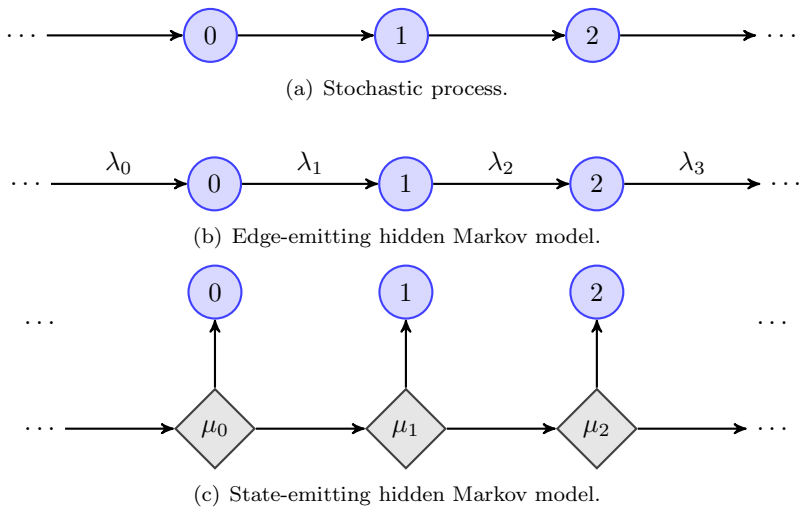

(a) Stochastic process.

(b) Edge-emitting hidden Markov model.

(c) State-emitting hidden Markov model.

FIGURE 9. A stochastic process and the two kinds of hidden Markov models.

7. Classical correlations as hidden Bayesian networks

Hidden Markov models. Hidden variables have been studied extensively not only in quantum foundations, but also in machine learning, and in particular hidden variables for stochastic processes. Recall that a **stochastic process** is a joint distribution $P(o[\mathbb{Z}])$ of a collection of random variables $o[\mathbb{Z}] = (o[t])_{t\in\mathbb{Z}}$ indexed by integer timesteps $t \in \mathbb{Z}$. The underlying causal structure is the one of $\mathbb{Z}$, as in Figure 9(a); the node labels are precisely the timesteps $t \in \mathbb{Z}$. We now informally discuss correlations and classical correlations on this causal structure; this must remain at an informal level since our definitions and results of the previous sections, like Definition 4.1, are not actually applicable to the case of infinite G, like $G = \mathbb{Z}$.

Informally speaking, every joint distribution $P(a_{\mathbb{Z}})$ should be a correlation in the sense of Definition 3.1, since there are no two non-empty subsets of $\mathbb{Z}$ with disjoint causal past. Realizations of $P(a_{\mathbb{Z}})$ in terms of classical hidden variables, as in Definition 4.1 and depicted in Figure 9(b), are known as **edge-emitting hidden Markov models** [28, Sec. 3.2][7]. Hidden Markov models have a wide range of applications e.g. in machine learning [46] and computational biology [47]. The essential idea is that the hidden variables $(\lambda_t)_{t\in\mathbb{Z}}$ form a Markov chain in the sense that λ_{t+1} may become correlated with λ_{t-1} only through λ_t, while the analogous statement may not be true for the observed values $o[t]$. Since Markov chains such as the hidden distribution $P(\lambda_{\mathbb{Z}})$ are much easier to handle mathematically and computationally, it is useful to consider a representation of $P(o[\mathbb{Z}])$, the actual stochastic process of interest, as a classical correlation. It is easy to see that every $P(o[\mathbb{Z}])$ actually admits such a representation by defining each hidden variable space (X_t, Σ_t) to be equal to $\prod_{t\in\mathbb{Z}} O_t$, so that one can take all hidden variables λ_t to be equal and to coincide with the entire sequence of outputs $(o[t])_{t\in\mathbb{Z}}$; an information processing gate $P(o[t]\lambda_{t+1}|\lambda_t)$ then simply

[7]There are other technical differences besides the fact that our definition only applies to finite graphs. The hidden variable spaces X_t in a hidden Markov model are typically assumed to be finite, and the stochastic maps $P(a_t\lambda_{t+t}|\lambda_t)$ are typically assumed to be independent of t. We use the term "hidden Markov model" in a sense which does not impose either of these requirements.

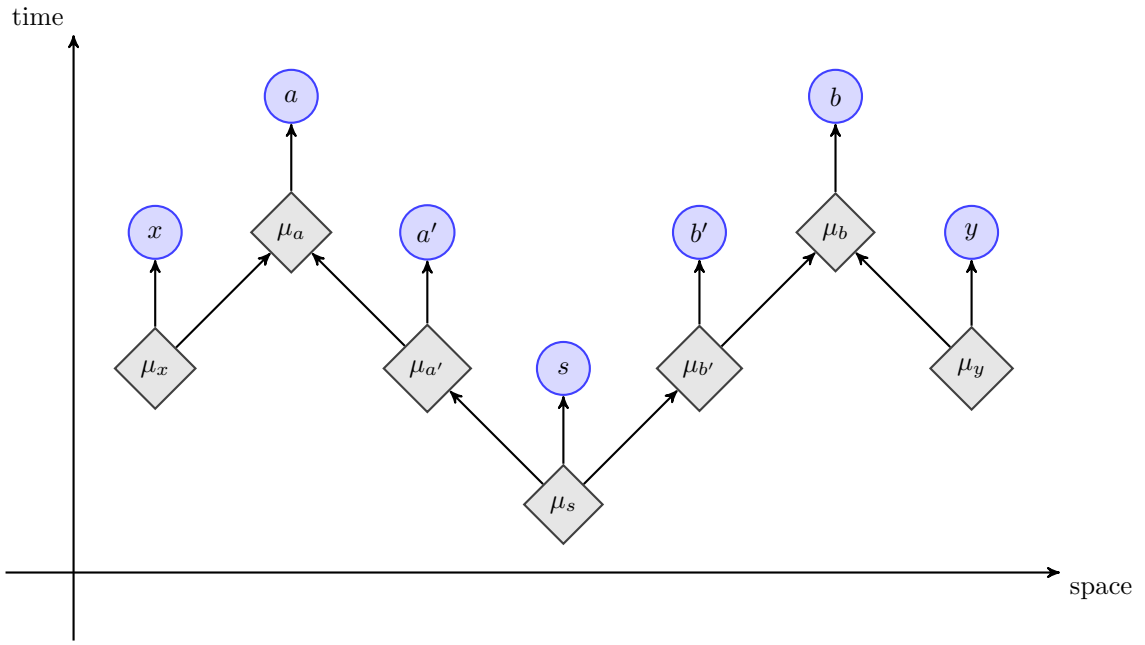

FIGURE 10. Hidden Bayesian network for Popescu's "hidden nonlocality" scenario (Figure 2).

takes an incoming $\lambda_t = o[\mathbb{Z}]$ and deterministically outputs the corresponding component $o[t]$ and forwards the entire sequence $o[\mathbb{Z}]$ as λ_{t+1}. On the other hand, in general it is *not* the case that every $P(o[\mathbb{Z}])$ can be represented as a classical correlation in such a way that all hidden variable spaces X_t are finite [**28**, Sec. 3.6].

So much for edge-emitting hidden Markov models. One more commonly encounters **state-emitting** hidden Markov models [**48**], of which Figure 9(c) shows an illustration. These are defined in a different way: again, there are hidden variable spaces (Y_t, Ω_t) carrying hidden variables $\mu_t \in Y_t$, but this time they come equipped with a pair of conditional distributions

$$P(\mu_{t+1}|\mu_t), \qquad P(o[t]|\mu_t),$$

which can now be thought of as living on the *nodes* in Figure 9(c). The first of these describes an updating of the hidden variable from μ_t to μ_{t+1}, while the second can be thought of as a probing of the hidden variable μ_t resulting in an outcome $o[t]$. The main point here is that these two updating operations are taken to be independent.

The purpose of this section is to generalize the latter notion of state-emitting hidden Markov model to an arbitrary (finite directed acyclic) graph $G = (V, E)$, which results in the notion of **hidden Bayesian network**. Then we prove that a correlation $P(o[V])$ is classical if and only if it admits a description in terms of a hidden Bayesian network. Figure 10 displays the structure of a hidden Bayesian network for the example causal structure of Figure 2.

7.1. DEFINITION. *A **Bayesian network** on a graph G is an assignment of a measurable space (Y_v, Ω_v) to every node $v \in V$ together with a conditional probability distribution*

$$P(\mu_v|\mu_{\mathrm{pa}(v)})$$

which generates the overall distribution

$$P(\mu_V) = \prod_{v \in V} P(\mu_v | \mu_{\mathrm{pa}(v)}). \tag{7.1}$$

Bayesian networks have been studied extensively for the case that the measurable spaces (Y_v, Ω_v) are finite [13], and for this case there are several equivalent conditions known for when a joint distribution $P(\mu_V)$ constitutes a Bayesian network.

7.2. Definition. *A correlation $P(o[V])$* ***arises from a hidden Bayesian network*** *if there is a Bayesian network with distribution $P(\mu_V)$ on G and stochastic maps $P(o[v]|\mu_v)$ for every $v \in V$ such that*

$$P(o[V]) = \int_{\mu_V} \left(\prod_{v \in V} P(o[v]|\mu_v) \right) \cdot P(\mu_V). \tag{7.2}$$

Intuitively, this formula expresses the idea that each $o[v]$ should be a noisy coarse-graining of the hidden variable μ_v, which are hidden variables whose dependencies are captured by the structure of the graph G.

7.3. Theorem. *A correlation $P(o[V])$ arises from a hidden Bayesian network if and only if it is classical.*

The proof will actually show that any joint distribution represented in the form (7.2) is a classical correlation; in other words, assuming that $P(o[V])$ is a correlation in Definition 7.2 is redundant.

Proof. We start with the "if" part, i.e. we assume a representation as a classical correlation in the form (4.2). We define the new hidden variable spaces (Y_v, Ω_v) to be given by, in terms of the old hidden variable spaces (X_e, Σ_e),

$$Y_v := O_v \times \prod_{e \in \mathrm{ch}(v)} X_e, \tag{7.3}$$

where it is understood that the σ-algebra Ω_v is given by the product of the corresponding Σ_e with $\mathbf{2}^{O_v}$. The basic idea here is that the new hidden variable $\mu_v \in Y_v$ will contain all the data which, in the representation as a classical correlation, corresponds to the output of the information processing gate $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$. The stochastic map $P(o[v]|\mu_v)$ which assigns to each new hidden variable a corresponding outcome can then simply be taken to be the projection from (7.3) onto O_v, and as such it is deterministic.

This identification $\mu_v = (o[v], \lambda_{\mathrm{out}(v)})$ of the new hidden variable μ_v with the outcome $o[v]$ together with the tuple of old hidden variables $\lambda_{\mathrm{out}(v)}$, which together form the output of the old information processing gate $P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)})$, suggests that we should also choose the conditional distributions in (7.1) to coincide with that information processing gate,

$$P(\mu_v|\mu_{\mathrm{pa}(v)}) := P(o[v]\lambda_{\mathrm{out}(v)}|\lambda_{\mathrm{in}(v)}).$$

This is a sensible definition if we consider $\lambda_{\mathrm{in}(v)}$ as a function of $\mu_{\mathrm{pa}(v)} = (o[\mathrm{pa}(v)], \lambda_{\mathrm{out}(\mathrm{pa}(v))})$, so that the left-hand side becomes a well-defined probability measure on (7.3) once a concrete tuple of values $\mu_{\mathrm{pa}(v)}$ has been specified.

With this definition of the hidden Bayesian network, we need to show that (7.2) indeed reproduces the original correlations. We have

$$P_{\text{new}}(o[V]) = \int_{\mu_V} \left(\prod_{v \in V} P(o[v]|\mu_v) \right) \cdot P(\mu_V) \overset{(7.1)}{=} \int_{\mu_V} \prod_{v \in V} P(o[v]|\mu_v)\, P(\mu_v|\mu_{\text{pa}(v)})$$
$$\overset{(*)}{=} \int_{\lambda_E} \prod_{v \in V} P(o[v]\lambda_{\text{out}(v)}|\lambda_{\text{in}(v)}) = P_{\text{old}}(o[V]),$$

where $(*)$ has used the fact that $P(o[v]|\mu_v)$ is simply given by projection onto the first factor, so that when decomposing μ_v into its components, only those tuples with the correct $o[v]$ component survive. Since every edge occurs in $\text{out}(v)$ for exactly one v, the resulting integral over μ_V turns into an integral over λ_E.

For the "only if" part, we start with a representation of the form (7.2). In this case, the new hidden variables are the λ_e, while the old ones are the μ_v. We define the new hidden variable spaces to be given by

$$X_e := Y_{\text{src}(e)}. \tag{7.4}$$

In other words, all outgoing edges $e \in \text{out}(v)$ at a given node v carry the same hidden variable space. Moreover, we construct the information processing gate at v so that all the outgoing information is the same on all these edges:

$$P(o[v]\lambda_{\text{out}(v)}|\lambda_{\text{in}(v)}) := P(o[v]|\mu_{\text{pa}(v)})\, P(\mu_v|\mu_{\text{pa}(v)}), \tag{7.5}$$

which is to be understood as follows: the known incoming values $\lambda_{\text{in}(v)}$ can be identified by (7.4) with values for the $\mu_{\text{pa}(v)}$. Then when keeping these values fixed, the right-hand side of (7.5) defines a probability measure on $O_v \times Y_v$ (actually a product measure). We can push this forward along the map $Y_v \to X_{\text{out}(v)}$, which by (7.4) we can take to be the inclusion along the diagonal. In other words, the resulting measure is supported on those tuples $\lambda_{\text{out}(v)}$ for which all components are the same. This ends the definition of the hidden variables and information processing gates.

We now verify that the resulting correlation is indeed the expected one,

$$P_{\text{new}}(o[V]) = \int_{\lambda_E} \prod_{v \in V} P(o[v]\lambda_{\text{out}(v)}|\lambda_{\text{in}(v)}) = \int_{\lambda_E} \prod_{v \in V} P(o[v]|\mu_{\text{pa}(v)})\, P(\mu_v|\mu_{\text{pa}(v)}),$$

where we simply plugged in (7.5), and it is understood that the intended interpretation of (7.5) needs to be applied. We can evaluate this further by noting that, at every node v, the integral over $\lambda_{\text{out}(v)}$ is supported on the diagonal, and hence it is sufficient to do an integral over μ_v. This proves that the previous expression is equal to

$$\int_{\mu_V} \prod_{v \in V} P(o[v]|\mu_{\text{pa}(v)})\, P(\mu_v|\mu_{\text{pa}(v)}) = P_{\text{old}}(o[V]),$$

as was to be shown. □

It should be interesting to figure out whether a similar statement can be derived for the quantum case, by using e.g. the definition of "quantum Bayesian network" of Leifer and Spekkens [49].

Appendix A. Measure-theoretic technicalities

Measurable spaces and measures. Here, we discuss the proper mathematics needed to describe classical hidden variables. A classical hidden variable is modelled by a set of possible values equipped with a probability distribution. Since we would like to allow this set of values to be of any size, finite or infinite, we first need to recall the standard mathematical structures used for dealing with probability distributions with an arbitrary number of outcomes. This is standard material can be found in any textbook on measure theory, e.g. [50]. The notions that we recall in the following can be thought of as classical analogues of concepts which are well-known to the quantum information and foundations community in the quantum-mechanical context.

A.1. Definition. *A **measurable space** is a pair (X, Σ) consisting of a set X and a collection of subsets $\Sigma \subseteq 2^X$ such that:*

- $\emptyset \in \Sigma$,
- *If $U \in \Sigma$, then also $X \backslash U \in \Sigma$,*
- *If $U_1, U_2, \ldots \in \Sigma$, then also $\bigcup_i U_i \in \Sigma$.*

If $U \in \Sigma$, then U is called a **measurable set**. The measurable sets U are those to which one expects to be able to assign a number measuring its "size", like the probability that a randomly chosen element of X will lie in U. A collection of sets Σ satisfying these axioms is called a **σ-algebra**.

We always assume that Σ *separates points* in X, which means that for any $x, y \in X$ there is a measurable set $A \in \Sigma$ such that $x \in A$, but $y \notin A$. If X is finite, this is equivalent to assuming that $\Sigma = 2^X$, meaning that all sets are actually measurable. The necessity to distinguish measurable from non-measurable sets arises only for infinite X.

The rôle of measurable spaces in classical probability theory is more or less analogous to the one of Hilbert spaces in quantum theory. The analogues of states are the measures:

A.2. Definition. *A **measure** on (X, Σ) is a function $P : \Sigma \to \mathbb{R}_{\geq 0}$ with*

- $P(\emptyset) = 0$,
- *$P\left(\bigcup_i U_i\right) = \sum_i P(U_i)$ for every sequence of pairwise disjoint measurable sets $U_1, U_2, \ldots \in \Sigma$.*

The **normalization** of a measure is the number $P(X)$. If $P(X) = 1$, then P is a **probability measure**; this models a situation in which it is certain that some outcome $x \in X$ occurs, and then we say that (X, Σ, P) is a **probability space**. Since we do not always want to make this assumption, we also allow that the total probability may be different from one, in which case it will typically be less than one. We nevertheless use the letter "P" to denote a measure since all our measures do in some way or the other have an interpretation as probabilities.

The rôle of an operator on a Hilbert space is played by the concept of measurable function:

A.3. Definition. *A **measurable function** on (X, Σ) is a function $f : X \to \mathbb{R}$ such that for all $a, b \in \mathbb{R}$, the set*

$$f^{-1}([a, b]) := \{x \in X \mid a \leq f(x) \leq b\}$$

is measurable, i.e. in Σ.

Measurable functions that are **bounded**, i.e. take values in a bounded subset of $\mathbb{R}$, can be integrated with respect to any measure P. This means that we can form the integral

$$\int_x f(x) P(x),$$

and this corresponds to the **expectation value** of f. Note that we have now written the integration variable formally as an argument of P. Throughout this paper, this should be understood to be symbolic notation for integral of a function f with respect to a measure P. It is symbolic since the measure $P(\{x\})$ of any one-element set $\{x\}$ is typically zero, and hence it is not meaningful in any sense to regard the measure P as a function of the variable x. For the precise definition of this integral we refer to the literature on measure theory [50].

Stochastic operations. In the main text, we also crucially need to deal with operations taking a collection of hidden variables as input and producing another collection of hidden variables as output. In order to define what exactly this is supposed to mean, we need to define **stochastic operations** between measurable spaces, which are the classical analogue of quantum operations. To this end, We are going to reproduce a variant of a definition of Panangaden [51]; the difference is that we allow measures with arbitrary normalization instead of only subnormalized ones, and we also speak of stochastic operations instead of "stochastic relations" in order to keep the terminology more in line with the standard jargon for the quantum case.

A.4. Definition. *Let (X, Σ_X) and (Y, Σ_Y) be measurable spaces. A **stochastic operation** from X to Y is a function*

$$P(y|x) \ : \ X \times \Sigma_Y \to [0, 1]$$

satisfying the following properties:

- *For every fixed $x \in X$, $P(y|x)$ is a measure on Y.*
- *For every fixed $U \in \Sigma_Y$, the function $P(U|x)$ is a bounded measurable function of x.*

The idea behind the notation "$P(y|x)$" is that for every $x \in X$, one obtains a measure on Y which, by abuse of notation, we also denote $P(y|x)$. If both X and Y are finite, $P(y|x)$ is simply a matrix of nonnegative numbers indexed by $x \in X$ and $y \in Y$. In any case, a stochastic operation, with our definition, has the property that the normalization of the measure $P(y|x)$ may depend on the value of x.

There is an important special case of a stochastic operation: if I is a one-element set, then the stochastic operations from I to Y are precisely the measures on Y. This is similar to how quantum states can themselves be regarded as maps from a trivial system to a non-trivial system [30].

A stochastic operation from X to Y can be composed with a stochastic operation from Y to Z to give a stochastic operation from X to Z:

$$P(z|x) := \int_y P(z|y)P(y|x). \tag{A.1}$$

In words, this means that for every fixed z, we integrate the function $y \mapsto P(z|y)$ with respect to the measure $P(y|x)$. We refer to [51] for the (non-trivial!) proof showing that this definition of composition of stochastic operations does indeed define a category, and we denote this category by $\mathtt{StOp}$.

There is an interesting subcategory $\mathtt{StOp}_1 \subseteq \mathtt{StOp}$ consisting of those stochastic operations $P(y|x)$ which are **normalized**, by which we mean that they satisfy

$$\int_y P(y|x) = 1$$

for all $x \in X$. For such a normalized stochastic operation, we also use the term **conditional distribution**.

If one only considers finite measurable spaces and normalized stochastic operations between those, one obtains a category which coincides with the category `FinStoch` discussed e.g. in [52].

Products of measure spaces and stochastic operations. For given measurable spaces (X, Σ_X) and (Y, Σ_Y), there is a way to form their **product space** $(X \times Y, \Sigma_X \otimes \Sigma_Y)$, where $X \times Y$ is the cartesian product. The σ-algebra $\Sigma_X \otimes \Sigma_Y$ is defined to contain all sets of the form $U \times V$ for $U \in \Sigma_X$ and $V \in \Sigma_Y$; but in order to obtain a measurable space, we need to close this collection of basic measurable sets by the operations of Definition A.1. See [50, Sec. 33] for more detail.

For stochastic operations $P(y_1|x_1)$ and $P(y_2|x_2)$, there also ought to exist a "product" operation corresponding to the idea of executing both $P(y_1|x_1)$ and $P(y_2|x_2)$ in parallel, analogous to the description of sequential execution described by (A.1). This product operation

$$P(y_1|x_1)\, P(y_2|x_2) \;:\; (X_1 \times X_2, \Sigma_{X_1} \otimes \Sigma_{X_2}) \longrightarrow (Y_1 \times Y_2, \Sigma_{Y_1} \otimes \Sigma_{Y_2})$$

is defined to map a pair of values (x_1, x_2) to the unique measure on $Y_1 \times Y_2$ having the property that it assigns to the set $U_1 \times U_2 \in \Sigma_{Y_1} \otimes \Sigma_{Y_2}$ the number $P(U_1|x_1) \cdot P(U_2|x_2)$; we refer to [50, Thm. 35.B] for the proof of uniqueness. We leave it to the reader to verify that this does indeed give a stochastic operation and end by noting that this product turns `StOp` into a **symmetric monoidal category**. In the particular case $X_1 = X_2 = I$, the product of stochastic operations specializes to a definition of product of measures.

Appendix B. Finite approximations of probability spaces

Here we prove some technical results which are of crucial importance for the proof of Theorem 4.4. We refer to [50] for the additional terminology and background from measure theory. For sets A and B, we write

$$A \triangle B := (A \setminus B) \cup (B \setminus A)$$

for the symmetric difference.

B.1. Lemma. *Let (X, Σ, P) be a probability space, $\Gamma \subseteq \Sigma$ an algebra of sets and $f : X \to S$ a measurable function to some finite set S. Then for every $\delta > 0$, there exist sets $B_s \in \Gamma$ indexed by $s \in S$ such that $P\left(B_s \triangle f^{-1}(s)\right) < \delta$, and the B_s partition X.*

The main point here is that $B_s \in \Gamma$: since f is measurable, we already know that $f^{-1}(s) \in \Sigma$, and the problem is to approximate these $f^{-1}(s)$ by elements of Γ.

Proof. We use induction on $|S|$. If $|S| = 1$, then there is nothing to prove, since we can simply take $B_s := X$. For general S, we choose any element $u \in S$. We then apply the approximation lemma of measure theory [50, 13.D] to obtain a set $B_u \in \Gamma$ with $P\left(B_u \triangle f^{-1}(u)\right) < \delta/2$.

The idea now is to restrict everything to $X \backslash B_u$ and apply the induction hypothesis to that. To do so, we define the function $f' : X \backslash B_u \to S \backslash \{u\}$ by taking $f'(x) := f(x)$ for all x with $f(x) \neq u$; we let f' map the other x's with $f(x) = u$ to an arbitrary element of $S \backslash \{u\}$. Then f' differs from the restriction of f to $X \backslash B_u$ on a set of measure less than $\delta/2$.

Applying the induction hypothesis to this situation, with δ replaced by $\delta/2$, gives a partition $X \backslash B_u = \bigcup_{s \in S \backslash \{u\}} B_s$ satisfying $P\left(B_s \triangle f'^{-1}(s)\right) < \delta/2$ for all $s \neq u$. Since we already know that $P\left(f'^{-1}(s) \triangle f^{-1}(s)\right) < \delta/2$ for all s, we conclude that

$$\begin{aligned} P\left(B_s \triangle f^{-1}(s)\right) &\leq P(B_s \triangle f'^{-1}(s)) + P(f'^{-1}(s) \triangle f^{-1}(s)) \\ &< \delta/2 + \delta/2 = \delta, \end{aligned}$$

as desired. □

B.2. Proposition. *Let $(X_k, \Sigma_k)_{k=1}^n$ be a family of measurable spaces equipped with a probability measure P on the product $\prod_k X_k$ and S a finite set. Then for every $\varepsilon > 0$ and any measurable function*

$$f \ : \ X_1 \times \ldots \times X_n \longrightarrow S$$

there exist a finite set S_k together with a measurable function $f_k : X_k \to S_k$ for each k, and a function

$$g \ : \ S_1 \times \ldots \times S_n \longrightarrow S$$

such that the diagram

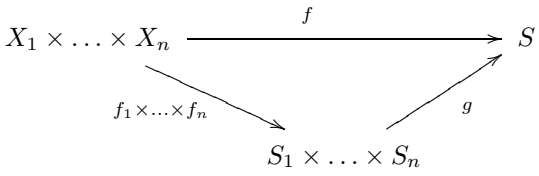

commutes up to ε, by which we mean that $f(x) = (g \circ (f_1 \times \ldots \times f_n))(x)$ holds with probability greater than $1 - \varepsilon$ with respect to P.

The main point here is that the S_k are finite: if we think of f as a function which coarse-grains $X_1 \times \ldots \times X_n$ down to a finite set S, then the claim is that this coarse-graining can be approximated arbitrarily well by first coarse-graining each X_k down to a finite set S_k, and then coarse-graining $S_1 \times \ldots \times S_n$ down to S.

Proof. The finite disjoint unions of products of measurable sets form an algebra of sets in $X_1 \times \ldots \times X_n$ [**50**, 33.E] which generates the product σ-algebra $\Sigma_1 \otimes \ldots \otimes \Sigma_n$. Hence Lemma B.1 gives a partition $X_1 \times \ldots \times X_n = \bigcup_{s \in S} B_s$ such that $P\left(B_s \triangle f^{-1}(s)\right) < \delta$. We will determine the required value of δ later. By construction, we know that each B_s is a finite union of sets of the form

$$C_1^i \times \ldots \times C_n^i, \tag{B.1}$$

where $C_k^i \in \Sigma_k$ and i indexes all the required sets for all s together. For fixed k, the finitely many sets C_k^i generate a finite Boolean algebra $\Omega_k \subseteq \Sigma_k$ whose atoms we denote by A_k^i. Again by construction, the B_s's are elements of $\Omega_1 \otimes \ldots \otimes \Omega_n$, and hence they can also be written as finite unions of products of atoms,

$$A_1^i \times \ldots \times A_n^i.$$

We then define S_k to be the set of all the atoms A_k^i, i.e. the spectrum of the Boolean algebra Ω_k. This S_k is guaranteed to be finite since Ω_k is finite. There is a canonical map $f_k : X_k \to S_k$ which sends an $x \in X_k$ to the unique atom which contains it.

We need to define the map $g : S_1 \times \ldots \times S_n \to S$. We do this by stipulating that a tuple of atoms $(A_1, \ldots, A_n) \in S_1 \times \ldots \times S_n$ maps to the unique $s \in S$ for which $A_1 \times \ldots \times A_n \subseteq B_s$; since the B_s's partition X and each of them is a finite union of products of atoms, that s is indeed unique.

Finally, we have to verify that this satisfies the required "commutativity up to ε" for a suitable choice of δ. This is easy to do once we realize that the composition $g \circ (f_1 \times \ldots \times f_n)$ is precisely the map which sends all of B_s to s. Hence the measure of the set of all tuples $(x_1, \ldots, x_n)$ on which f differs from $g \circ (f_1 \times \ldots \times f_n)$ is given by

$$P\left(\bigcup_{s \in S} (B_s \triangle f^{-1}(s))\right) \leq \sum_{s \in S} P\left(B_s \triangle f^{-1}(s)\right) < |S| \cdot \delta.$$

So if we do the whole construction with $\delta \leq \varepsilon / |S|$, the claim follows. □

References

[1] John S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964. ↑2, 9, 31

[2] Nicolas Brunner, Daniel Cavalcanti, Stefano Pironio, Valerio Scarani, and Stephanie Wehner. Bell nonlocality, 2014. ↑2, 5, 9, 15, 16, 31

[3] Artur K. Ekert. Quantum cryptography based on Bell's theorem. *Phys. Rev. Lett.*, 67:661–663, Aug 1991. ↑2

[4] Umesh Vazirani and Thomas Vidick. Fully device independent quantum key distribution. *Phys. Rev. Lett.*, 113:140501, 2014. ↑2, 11

[5] S. Pironio, A. Acín, S. Massar, A. Boyer de la Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe. Random numbers certified by Bells theorem. *Nature*, 464:1021, 2010. ↑2

[6] Umesh Vazirani and Thomas Vidick. Certifiable quantum dice: or, true random number generation secure against quantum adversaries. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 61–76, 2012. ↑2, 11

[7] George Svetlichny. Distinguishing three-body from two-body nonseparability by a Bell-type inequality. *Phys. Rev. D*, 35:3066, 1987. ↑2

[8] Sandu Popescu. Bell's inequalities and density matrices: Revealing "hidden" nonlocality. *Phys. Rev. Lett.*, 74:2619–2622, Apr 1995. ↑2, 3, 35

[9] Cyril Branciard, Nicolas Gisin, and Stefano Pironio. Characterizing the nonlocal correlations created via entanglement swapping. *Phys. Rev. Lett.*, 104:170401, 2010. ↑2, 4, 35

[10] Tobias Fritz. Beyond Bell's theorem I: correlation scenarios. *New J. Phys.*, 14:103001, 2012. ↑2, 4, 5, 7, 16, 35

[11] Rodrigo Gallego, Lars Erik Würflinger, Rafael Chaves, Antonio Acín, and Miguel Navascués. Nonlocality in sequential correlation scenarios. *New J. Phys.*, 16:033037, 2014. ↑2, 6, 35

[12] Lucien Hardy. The operator tensor formulation of quantum theory. *Phil. Trans. R. Soc. A*, 28:3385–3417, 2012. ↑3, 23, 31

[13] Judea Pearl. *Causality*. Cambridge University Press, Cambridge, second edition, 2009. Models, reasoning, and inference. ↑3, 4, 8, 9, 13, 38

[14] Albert Einstein, Boris Podolsky, and Nathan Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, 1935. ↑4

[15] Christopher J. Wood and Robert W. Spekkens. The lesson of causal discovery algorithms for quantum correlations: Causal explanations of Bell-inequality violations require fine-tuning. *New. J. Phys.*, 17:033002, 2015. ↑4, 31

[16] Joe Henson, Raymond Lal, and Matthew F. Pusey. Theory-independent limits on correlations from generalised bayesian networks. *New J. Phys.*, 16:113043, 2014. ↑4, 7, 16, 22, 29, 35

[17] Rafael Chaves, Lukas Luft, and David Gross. Causal structures from entropic information: geometry and novel scenarios. *New. J. Phys.*, 16:043001, 2014. ↑4, 16

[18] Rafael Chaves, Lukas Luft, Thiago O. Maciel, David Gross, Dominik Janzing, and Bernhard Schölkopf. Inferring latent structures via information inequalities. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, pages 112–121. AUAI Press, 2014. ↑4

[19] Rafael Chaves, Christian Majenz, and David Gross. Information-theoretic implications of quantum causal structures. *Nat. Comm.*, 6:5766, 2015. ↑4

[20] Rafael Chaves, Richard Kueng, Jonatan Bohr Brask, and David Gross. A unifying framework for relaxations of the causal assumptions in bell's theorem. *Phys. Rev. Lett.*, 114:140403, 2015. ↑4

[21] Frank Arntzenius. Reichenbach's common cause principle. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, 2010. plato.stanford.edu/archives/fall2010/entries/physics-Rpcc/. ↑5

[22] Ognyan Oreshkov, Fabio Costa, and Časlav Brukner. Quantum correlations with no causal order. *Nat. Comm.*, 3:1092, 2012. ↑7

[23] Peter Bierhorst. A rigorous analysis of the Clauser-Horne-Shimony-Holt inequality experiment when trials need not be independent. *Found. Phys.*, 44:736–761, 2014. ↑7

[24] Joe Henson. The causal set approach to quantum gravity, 2006. arXiv:gr-qc/0601121. ↑9

[25] Jonathan Barrett, Noah Linden, Serge Massar, Stefano Pironio, Sandu Popescu, and David Roberts. Nonlocal correlations as an information-theoretic resource. *Phys. Rev. A*, 71:022101, Feb 2005. ↑12

[26] Anthony J. Leggett and Anupam Garg. Quantum mechanics versus macroscopic realism: is the flux there when nobody looks? *Phys. Rev. Lett.*, 54:857, 1985. ↑13

[27] Arthur Fine. Hidden variables, joint probability, and the Bell inequalities. *Phys. Rev. Lett.*, 48(5):291–295, Feb 1982. ↑16, 19

[28] Daniel Ray Upper. *Theory and Algorithms for Hidden Markov Models and Generalized Hidden Markov Models*. PhD thesis, University of California at Berkeley, 1989. Available at csc.ucdavis.edu/∼cmg/papers/TAHMMGHMM.pdf.gz. ↑17, 36, 37
[29] Giulio Chiribella, G. Mauro D'Ariano, and Paolo Perinotti. Probabilistic theories with purification. *Phys. Rev. A*, 81:062348, 2010. ↑22, 23
[30] Bob Coecke and Eric Oliver Paquette. Categories for the practising physicist. In *New Structures for Physics*, pages 173–286. Springer, 2011. ↑22, 41
[31] David Spivak. Category theory for scientists (old version), 2013. arXiv:1302.6946. ↑22
[32] Bob Coecke and Raymond Lal. Causal categories: relativistically interacting processes. *Found. Phys.*, 43(4):458–501, 2012. ↑23
[33] Bob Coecke. Quantum picturalism. *Contemporary Physics*, 51(1):59–83, 2010. ↑23
[34] André Joyal and Ross Street. The geometry of tensor calculus, I. *Adv. Math.*, 88:55–112, 1991. ↑23, 24, 25, 28
[35] Tom Leinster. Higher operads, higher categories, 2003. arXiv:math/0305049. ↑24
[36] Bob Coecke. Terminality implies non-signalling. In *Proceedings 11th workshop on Quantum Physics and Logic*, volume 172, pages 27–35, 2014. ↑25
[37] David I. Spivak, Patrick Schultz, and Dylan Rupel. String diagrams for traced and compact categories are oriented 1-cobordisms, 2015. arXiv:1508.01069. ↑28
[38] Jonathan Barrett. Information processing in generalized probabilistic theories. *Phys. Rev. A*, 75(3):032304, Mar 2007. ↑29
[39] Anthony Short and Jonathan Barrett. Strong nonlocality: A trade-off between states and measurements. *New J. Phys.*, 12:033034, 2010. ↑29
[40] Nicolas Gisin. Non-realism: deep thought or soft option? *Found. Phys.*, 42:80–85, 2012. ↑31
[41] Nicolas Gisin. A possible definition of a *realistic* physics theory, 2014. arXiv:1401.0419. ↑31
[42] Paul R. Halmos. *A Hilbert Space Problem Book*, volume 19 of *Graduate Texts in Mathematics*. Springer, second edition, 1982. ↑31
[43] Richard V. Kadison and John R. Ringrose. *Fundamentals of the theory of operator algebras. Vol. I*, volume 100 of *Pure and Applied Mathematics*. Academic Press Inc. Harcourt Brace Jovanovich Publishers, 1983. Elementary theory. ↑31
[44] John B. Conway. *A Course in Operator Theory*, volume 21 of *Graduate Studies in Mathematics*. American Mathematical Society, 1999. ↑31
[45] Michael Nielsen and Isaac Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. ↑31
[46] Mark Gales and Steve Young. The application of hidden Markov models in speech recognition. *Foundations and Trends in Signal Processing*, 1(3):195–304, 2008. ↑36
[47] Sean R. Eddy. What is a hidden markov model? *Nature Biotechnology*, 22:1315–1316, 2004. ↑36
[48] Zoubin Ghahramani. An introduction to hidden Markov models and Bayesian networks. *Int. J. Patt. Recogn. Artif. Intell.*, 15(1):9–42, 2001. ↑37
[49] Matthew Leifer and Robert Spekkens. Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Phys. Rev. A*, 88:052130, 2013. ↑39
[50] Paul R. Halmos. *Measure theory*. D. Van Nostrand Company, Inc., New York, N. Y., 1950. ↑40, 41, 42, 43, 44
[51] Prakash Panangaden. *Labelled Markov Processes*. Imperial College Press, 2009. ↑41
[52] John C. Baez and Tobias Fritz. A Bayesian characterization of relative entropy. *Theory and Applications of Categories*, 29:421–456, 2014. ↑42

Tobias Fritz, Perimeter Institute for Theoretical Physics, Waterloo ON, Canada
E-mail address: tfritz@perimeterinstitute.ca